Featured in Physics

Continual Evolution in Nonreciprocal Ecological Models

Aditya Mahadevan
Department of Physics, Stanford University, Stanford, California 94305, USA

Daniel S. Fisher
Department of Applied Physics, Stanford University, Stanford, California 94305, USA

(Received 3 December 2024; accepted 18 June 2025; published 4 August 2025)

Feedbacks between evolution and ecology are ubiquitous, with ecological interactions determining which mutants are successful, and these mutants in turn modifying community structure. We study the evolutionary dynamics of several ecological models with overlapping niches, including consumer resource and Lotka-Volterra models. Evolution is assumed slow, with ecological dynamics producing a stable community (with some strains going permanently extinct) between successive introductions of an invader or mutant. When new strains are slowly added to the community, either from an external pool or by large-effect mutations of exant strains, an initial evolutionary transient gives way to a diverse eco-evolutionary steady state. In this "Red Queen" phase of continual evolution, the biodiversity continues to turn over without the invasion probability of new variants getting smaller. For resource-mediated interactions, the Red Queen phase obtains for any amount of asymmetry in the interactions between strains, and is robust to "general fitness" differences in the intrinsic growth rates of strains. Via a dynamical mean-field theory framework valid for high-dimensional phenotype space, we analytically characterize the Red Queen eco-evolutionary steady state in a particular limit of model parameters. Scaling arguments enable a more general understanding of the steady state and evolutionary transients toward it. This work therefore establishes simple models of continual evolution in an ecological context *without* host-pathogen arms races and points to the generality of Red Queen evolution. However, we also find other eco-evolutionary phases in simple models: For generalized Lotka-Volterra models with near-symmetric interactions an "oligarch" phase emerges in which the evolutionary dynamics continually slow down and a substantial fraction of the community's abundance condenses into a handful of slowly turning-over strains.

DOI: 10.1103/dv3k-75b9

I. INTRODUCTION

Ecological and evolutionary processes are inextricably linked. In diverse microbial communities with large population sizes, there is constant interplay between ecological interactions and variation generated by mutation. Evolution continually creates new variants that interact with the community differently from their parent strains, and these ecological interactions determine the community's properties and the fate of the new mutants, thereby influencing future evolution. As the increasing body of data from natural microbial communities shows, microbes live in close proximity with one another, with multiple strains and substrains of the same species coexisting in many natural environments [1–3]. Understanding how ecological interaction webs influence evolution—and vice versa—is crucial for developing a richer picture of microbial evolution in both experimental and natural settings. A major puzzle is the coexistence of closely related microbial strains, as evinced by the broad distribution of genomic divergences found within bacterial species apparently occupying the same ecological niche [4,5]. How does such diversity evolve and persist without clear ecological differences stabilizing the coexistence of strains? Theory is needed to develop distinguishable scenarios for the evolution and maintenance of fine-scale biodiversity in large microbial populations.

Even in controlled laboratory experiments with initially clonal microbial populations, the development of ecology is inevitable. Although *apparent* selection pressures in the laboratory can be held constant, evolution in "simple" environments still generically results in the diversification and coexistence of different strains [6–8]. Long after increases in fitness (as measured by competition with the ancestor) have slowed and populations seem to be well adapted to laboratory conditions, it has been observed (for example in the long-term *E. coli* evolutions of Richard Lenski) that the rate of successful mutations does not continually decrease [9]. Furthermore, diversification and coexistence of lineages is seen even in such experiments *designed* to be simple enough to eliminate ecological effects [9,10]. Remarkably, even in laboratory experiments on bacterial populations without an external nutrient supply which might be expected to exhibit severe resource limitation and diversity collapse—diversification and phenotypic turnover is found to continue over a timescale of years [11].

With interacting bacteria and phage populations, there is clear evidence—both experimental and observational—for

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the continual turnover and diversification of microbial lineages [12–14]. The "Red Queen hypothesis" [15] is often invoked to describe such arms-race evolutionary dynamics of hosts and pathogens: as its eponymous character tells Alice in Lewis Carroll's *Through the Looking Glass*, "it takes all the running you can do to keep in the same place" [16]. In the current work we use "Red Queen dynamics" more generally to describe continual turnover of biodiversity in an evolving community, without successful invasions of new mutants becoming rarer. We look for such dynamics in ecological contexts other than that of hosts and pathogens.

Inspired by observations of microbial populations, we aim to understand what generic scenarios can emerge in the long-term dynamics of evolution of ecological communities. We are particularly interested in robust behaviors that obtain across a range of parameters and models—such "phases" of eco-evolutionary dynamics present a set of qualitative possibilities that one can hope to compare with experiment or observation. The phases that we study emerge in the limit of large biological complexity: high-dimensional phenotype spaces in which evolutionary and ecological processes play out. Looking for phases that are robust in the limit of high dimensionality D—roughly the number of phenotypic properties of organisms that affect the ecological dynamics—increases the likelihood that the conclusions we draw are somewhat universal and relevant to natural systems, which are more complex—in multiple ways—than the simple models we study.

To understand what long-term behaviors emerge in high-dimensional eco-evolutionary models, it is useful to first think separately about the dynamics of evolution and ecology, initially in a well-mixed population with *no spatial structure*. A further simplification is natural for large microbial populations: that the dynamics are essentially deterministic, with demographic fluctuations negligible except close to extinctions.

A. Evolution without ecology

In a static environment without ecology, gradual evolution of a large almost-clonal population can be caricatured as gradient ascent on a complex fitness landscape. The fitness landscape is a map from some high-dimensional organismal phenotype to the "fitness": a misleading but canonical term for a measure of reproductive success that is neither a simple scalar nor a useful quantity except in the simplest nonecological contexts. The emergence and fixation of beneficial mutations drives the population to high-fitness regions of phenotype space and evolution slows down as the population approaches one of the (typically) many local maxima of the fitness landscape [17]. However, in high dimensions, this simple intuition is only part of the story: the phenotype will generally spend a long time wandering around saddles before "committing" to a particular maximum, which is likely to be only marginally stable [18].

In large populations with an abundant supply of mutations, the evolutionary dynamics of the population of genotypes is complex, with simultaneous exploration of multiple phenotypic directions: this regime has not been much studied theoretically. However, if there is no ecology except resource limitation capping the total population size, eventually the behavior is simpler: at long times the supply of beneficial mutations will become so limited that the population will again become almost clonal. By contrast, in an extrinsically changing environment, the fitness landscape depends on time, and is better caricatured as a "seascape." Not surprisingly, a continually changing environment can lead to evolutionary dynamics which proceed indefinitely without slowing down [19] and produce continually turning-over diversity: again, such phenomena have not been much studied for large populations in high phenotypic dimensions.

B. Ecology without evolution

Just as evolutionary dynamics of large populations can display rich behavior even without ecology, ecological dynamics with no evolution can exhibit various behaviors, with globally stable, multistable, and chaotic states all possible in simple models [20–22]. The most-studied contexts are ecological dynamics of *top-down assembled communities*, which consist of a large number of species brought together with their abundances obeying some prescribed ecological dynamics. The simplest ecological phase that obtains at long times is a state in which species abundances reach a globally attracting stable fixed point with some number of species (the stable community) having nonzero abundance and the rest going extinct, unable to reinvade even if reintroduced. As the diversity of this assembled community increases, the stable phase can yield to a more complicated behavior exhibiting multiple attractors and transient chaos, which remains only partially understood [20,22,23]. Chaotic fluctuations in the high-diversity unstable regime lead to many extinctions and the system settles down to one of many possible stable communities or limit cycles, but these fixed points are generically unstable to invasion by some of the extinct strains. A population of dormant spores that germinates and resuscitates extinct populations can sustain the chaos for times exponentially long in the lifetimes of the spores, but as a matter of principle such dynamics do not constitute a true ecological phase [24].

C. Interplay of ecology and evolution

To bring evolution and ecology together, we focus here on the slowly evolving limit in which the ecological dynamics are much faster than the evolutionary dynamics and spatial mixing is fast enough to be ignored. This regime is the hardest in which to find sustained diversity and continual evolution because it does not allow short-term transient coexistence driven by continual invasions of new mutants, which was studied in Refs. [25–27]. Motivated especially by the puzzle of fine-scale diversity, we are primarily interested in multiple strains of a single species with overlapping niches.

We study models in which the ecological dynamics are simplest, converging to a (usually) globally attracting fixed point with some strains permanently extinct. Under slow evolution, do the resulting eco-evolutionary dynamics remain diverse and continually turn over in a "Red Queen" steady state, or does evolution get progressively slower and slower? Do generalist mutants take over and cause the diversity to crash? What phases of eco-evolutionary dynamics exist as a function of model parameters?

It has been shown that Red Queen dynamics emerge generically in models of an almost-clonal population evolving in a high-dimensional fitness landscape with a small amount of environmental feedback [18]. In these models, the fitness landscape—now better caricatured as a "snowscape"—changes as the evolving population affects its environment. For even a small amount of such feedback, the population wanders through phenotype space without settling down to a stable fixed point, due to selective forces that arise from the gradient of its continually modified fitness snowscape. Universality emerges in these models in the limit of large phenotype dimension, with different microscopic models displaying similar phase diagrams, including a phase of sustained evolutionary chaos [18].

In the eco-evolutionary models we study here, ecology acts loosely as a manifestation of feedback between organism and environment, with the "environment" of any particular strain set by its interactions with the rest of the community, and thus subject to change every time a new mutant arises in the community. Can natural models of this process result in a Red Queen phase of evolution? If so, what sets the diversity of the community in this phase?

Some previous work has found that ecological communities, although capable of sustaining extensive diversity through standard mechanisms of resource competition [28] lose much of their diversity when they start to gradually evolve either by small mutations or by the introduction of independent species [29,30]. These results were argued to deepen the puzzle of the extensive diversity observed in nature. However, the models in which such results have been reported are complicated enough that analytical understanding is difficult, and generality of the claims is not clear—indeed other work has suggested that phenotypic complexity *allows* the gradual evolution of high diversity [31].

D. Eco-evolutionary dynamics

The evolutionary process we study involves adding strains one by one to an ecological community, with extinctions permanent when they occur. Once extinct, a strain will never be able to re-emerge, even if conditions become favorable for its growth. Therefore the order in which strains are introduced influences the fate of the community. However, the *statistical* properties of the evolving community for large phenotypic dimension D depend only on the pool from which invaders are drawn, rather than on the specific properties of the invaders.

Previous work has studied the process of adding strains into ecological models *without* permanent extinctions [32–35], where re-invasion of previously extinct strains is possible if conditions later become favorable for them. This process eliminates the important difference between communities assembled all at once (top-down) and gradually (bottom-up): the latter is our focus.

A large amount of work by theoretical ecologists and statistical physicists has focused on ecological models that are nongeneric because they admit a Lyapunov function which is always increasing under the eco-evolutionary dynamics [35–39]. The existence of such a Lyapunov function depends on the symmetry of ecological interactions [40], e.g., through perfect efficiency in the conversion of nutrients to biomass, which does not hold in biological systems. Therefore, the generality of conclusions drawn from such models is unclear. In the current work we focus on generic behavior, drawing conclusions away from the special slice of parameter space with a Lyapunov function. Indeed we show that, under evolution, the behavior of models with a Lyapunov function is qualitatively different from the general case. Several other studies [41–43] have noted the importance of *nonreciprocal* or *asymmetric* ecological interactions for reaching a state of continual evolution—we provide further evidence for this conclusion in the high-diversity limit, which complements the more often-studied low-diversity case [41,42].

E. Models

We are interested in dynamics of closely related strains, where differences between strain interactions with one another are sums of many positive and negative effects. To caricature this biological complexity, we parametrize the strain-to-strain differences in phenotypic properties as *random variables* drawn from some distribution. This approximation of unknown complexity by randomness is a useful way to get a sense of general patterns of behavior that might apply across a range of ecological-evolutionary situations [44].

In the simplest models, strains interact with one another *directly*, with all pairwise interactions chosen from some specified ensemble. These models, which include generalized Lotka-Volterra interactions [20], suffer from a lack of strain phenotypes which should dictate how each strain interacts with any other extant or prospective strain. As a result, the effective phenotypic dimensionality in Lotka-Volterra models is infinite. Problematically, such models do not generally admit a straightforward way of drawing random mutants correlated with a parent strain in the community [45]. Nevertheless, we find that the limit of infinite phenotype dimension (with invading strains drawn independently from the extant community) is instructive for understanding evolution driven by mutant strains correlated with their parents.

To explicitly study the effects of parent-mutant correlations we introduce structure into the ecological interactions, considering the case where they are mediated by a collection of chemicals (or resources) in the environment. Each strain has a phenotype which describes its effect on the resources and the resources' effect on its growth rate—these phenotypes consequently determine the interstrain interactions. Small mutations correspond to having highly correlated parent and mutant phenotypes, with the amount of similarity between mutant and parent playing an important role.

The phenotype of each strain in such *consumer resource* models is given by its growth and consumption rates for each of the resources. The number of resources—or more generally, interaction-mediating chemicals—is the key quantity D, which sets the dimensionality of the environment of the effective consumer-consumer interactions. It also sets a natural scale for the number of strains that can coexist: in the models we study here, the diversity is bounded above by D, and "high diversity" means that the number of coexisting strains is $O(D)$. The limit of large D—surely biologically relevant [31]—will allow us to analyze statistical (and hopefully universal) properties of the resulting eco-evolutionary dynamics.

When mutational effects are sufficiently large, we find that key aspects of the evolutionary dynamics are captured by the dynamics of *bottom-up assembly* (with extinctions), wherein phenotypically *independent* invaders enter the community serially. The understanding from this simplified regime gives us qualitative insight into differences between a top-down-assembled and an evolved community, along with a framework for understanding the Red Queen phase that occurs for correlated mutants.

F. Absence of strict metabolic tradeoffs

In order for the Red Queen phase to be robust, it must be stable to the presence of "general fitness" mutations which increase the growth rate of a strain *regardless* of the composition of the community. These general fitnesses appear naturally in models of resource competition, where they correspond to the growth rates of strains in isolation. Metabolic trade-offs, which constrain the width of the general fitness distribution, are sometimes invoked to sustain high levels of diversity in ecological settings. These trade-offs can be strict [46,47] or—more realistically—allow some variation in the total metabolic budget [38]. Without strict trade-offs, can consumer resource and Lotka-Volterra models evolve towards a diverse state in which the distribution of extant general fitnesses is narrow? Or does evolution constantly push into the tail of the general fitness distribution and fail to reach a steady state? Do generalist species emerge and reduce diversity by outcompeting other strains?

II. MAIN RESULTS AND RELATED WORKS

Our main result is the existence of a robust high-diversity Red Queen phase of eco-evolutionary dynamics across a number of models. This phase is characterized by continual evolution and turnover of biodiversity without systematic improvement of any kind: therefore the probability of a randomly drawn external strain invading the community does not get arbitrarily small. The Red Queen phase occurs across a range of parameters in various models provided there is sufficient nonreciprocity in the ecological interactions, and it persists in the presence of general fitness differences, as long as their distribution decays fast enough. The tail of the general fitness distribution controls the probability of successful invasion and therefore the rate of turnover in the steady state, but if the tail decays fast enough it does not limit the diversity. We are able to analytically understand the Red Queen phase through a simple exactly solvable limit of infinite phenotypic dimensionality, and we make scaling predictions that are borne out in simulations across a range of parameters.

Despite the generality of Red Queen eco-evolutionary phase, we find that it does not always occur. In a class of related models, as the ecological interactions become more reciprocal, we see a qualitatively distinct possibility for long-term evolution which we term the "oligarch" phase—a counterpart to the Red Queen's monarchy—in which a small number of strains dominate the total abundance of the community and turn over increasingly slowly, with the invasion probability of a random outsider continually decreasing. We have not understood the oligarch phase quantitatively, but its presence is an important feature that shows the possibility of different eco-evolutionary phases even in simple models.

To put our analysis in the context of previous work, we note that Ref. [35] studied the effect of a single small mutation on a diverse community, but did not address the long-term outcome of such an evolutionary process and did not consider permanent extinctions or asymmetric interactions. Reference [48] studied long-term evolution but mostly analyzed the case of a small number of substitutable resources. Reference [49] studied bottom-up assembly with extinctions but did not explore parameter space and crucially neglected the effect of a distribution of general fitnesses. Reference [50] studied a different regime of a related *ecological* model which displays aspects of the phenomena we observe, and the authors used similar analytical tools to those we employ. We later comment further on the relation between our work and that of Ref. [50].

III. CONSUMER RESOURCE MODEL

A. Ecological dynamics and stable communities

We first study a standard model of resource competition with D externally supplied resources (or more generally, interaction-mediating chemicals), and an assembled population of S different strains that consume these resources [37,38,40,46]. (See Table I for notations.) The ecological dynamics of this consumer resource (CR) model are

$$\frac{dn_i}{dt} = n_i\left(\sum_{\beta=1}^{D}\mathcal{G}_{i\beta}\mathcal{R}_\beta - d\right), \tag{1a}$$

$$\frac{d\mathcal{R}_\alpha}{dt} = K_\alpha - \mathcal{R}_\alpha\sum_{j=1}^{S}\mathcal{F}_{j\alpha}n_j - \omega\mathcal{R}_\alpha, \tag{1b}$$

where the $\{n_i\}$ are the populations of strains, labeled by Roman indices $i \in \{1, 2, \ldots S\}$ and the $\{\mathcal{R}_\alpha\}$ are resource abundances labeled by Greek indices $\alpha \in \{1, 2, \ldots D\}$. The resource supply vector $\mathbf{K}$ has components K_α, and d and ω are respectively the death rate of the consumers and the decay (or dilution) rate of the resources. Without substantial loss of generality, we take all the K_α to be equal and given by a constant K, and we set $\omega = 0$ to reduce the number of parameters without changing the qualitative behavior. The $S \times D$ matrices $\mathcal{G}$ and $\mathcal{F}$ respectively parametrize the *growth* and *feeding* rates of each consumer on each resource: the phenotype of strain i is the corresponding row of $\mathcal{G}$ and $\mathcal{F}$ taken together, so that the dimensionality of phenotype space is $2D$. The elements of $\mathcal{G}$ and $\mathcal{F}$ have some positive mean—an average growth or feeding rate—in addition to a heterogeneous part that varies between strains and resources.

Reference [38] studied the ecological dynamics of a *symmetric* version of this model with $\mathcal{F} = \mathcal{G}$. Although we expect $\mathcal{G}$ and $\mathcal{F}$ to be positively correlated, some amount of asymmetry—arising from, e.g., varying efficiencies in utilization of resources—is generic. We show that this asymmetry is important for the long-term evolutionary dynamics.

TABLE I. Definitions of variables and terms.

CR	Consumer resource model [Eq. (1)]
LV	Lotka-Volterra model [Eq. (F1)]
S	Number of strains in initial pool
D	Phenotypic dimension; number of resources
L	Number of extant strains at fixed point
n_i	Abundance of strain i
$\mathbf{K}$	Vector of resource supply rates
U_i	General fitness in CR model; $U_i = \sum_\alpha \mathcal{G}_{i\alpha} K_\alpha$
$\mathcal{G}$	Growth rate matrix
$\mathcal{F}$	Feeding rate matrix
κ	Correlation between growth and feeding; $\kappa \in [0, 1]$
ξ_i	Invasion fitness = bias
ν_i	Fractional abundance of strain i; $\sum_i \nu_i = 1$
ζ_i	Drive = community-dependent component of bias
σ_ζ	Width of cavity drive distribution
s_i	General fitness; intrinsic component of bias
$p(s)$	Distribution of the s_i
Σ	Scale of s distribution
Υ	Community mean fitness
ψ	$p(s)$ decays as $\exp[-(s/\Sigma)^\psi/\psi]$
$\hat{A}$	Value of A in the saturated assembled community
$\overline{A}$	Evolutionary time average of A
G	Matrix of growth rate differences; mean 0
F	Matrix of feeding rate differences; mean 0
$\mathcal{V}$	Pairwise strain interaction matrix $= -GF^\intercal$
V	Scaled interaction matrix; $V = \mathcal{V}/\sqrt{D}$
γ	Symmetry parameter for LV interactions; $\gamma \in [-1, 1]$
Q	Niche parameter of LV model; $\mathrm{E}[V_{ii}] = -Q$
p_{inv}	Probability of a new variant successfully invading
ρ	Correlation between parent and mutant
$\Delta(\Upsilon)$	Scale of s distribution around $s \approx \Upsilon$ [Eq. (14)]

1. Interaction ensemble

For closely related strains, the differences between interactions will be sums of many positive and negative effects: thus we treat $\mathcal{G}_{j\alpha}$ and $\mathcal{F}_{j\alpha}$ as random variables drawn from some distribution. With the interpretation of consumption and growth, one expects the elements of $\mathcal{F}$ and $\mathcal{G}$ to be positive, although a small fraction being negative (corresponding loosely to cross feeding or inhibitory effects) does not change the behavior. For simplicity we draw the elements from Gaussian distributions with means μ_g, μ_f and variances σ_g^2, σ_f^2, respectively. The choice of a Gaussian distribution is convenient as we later take advantage of this structure to evolve strains by drawing from the ensemble of interactions conditional on inequality constraints. However, as we show in Sec. V D, the upper tails of the distributions of the elements of $\mathcal{G}$ and $\mathcal{F}$ play an important role in the evolution.

We parametrize elementwise correlations between the Gaussian entries of $\mathcal{G}$ and $\mathcal{F}$ by $\mathrm{Corr}[\mathcal{G}_{i\alpha}\mathcal{F}_{j\beta}] = \kappa\delta_{ij}\delta_{\alpha\beta}$, where $\mathrm{Corr}[A, B]$ is the Pearson correlation coefficient between A and B, and the *symmetry parameter* $\kappa \in [0, 1]$ describes the similarity between growth and consumption rates. With the choice of $\omega = 0$ there are only three important parameters other than S and D: κ, σ_g/μ_g, and σ_f/μ_f (Appendix A). The ratios σ_g/μ_g and σ_f/μ_f determine what fraction of the entries of $\mathcal{G}$ and $\mathcal{F}$ are positive; we consider $\sigma_f = \mu_f/4$ and $\sigma_g = \mu_g/4$ so that most entries are positive. As K only sets a scale for the consumer abundances, without loss of generality we set $\mu_f = \mu_g = 1$, $d = 1$, and $K = 1$.

The degree of correlation, κ, between the feeding and growth rates plays an important role: $\kappa = 1$ presumes identical efficiencies of conversion between resources and biomass between strains [38]. This symmetry guarantees the existence of a Lyapunov function (see Appendix A) that is always increasing along the ecological (and later evolutionary) dynamics of the CR model [37]. However, for generic $\kappa < 1$ the dynamics crucially do not have a Lyapunov function.

2. Ecological fixed points

One must supplement the deterministic dynamics of Eq. (1) with a small abundance threshold below which strains go *extinct*: their abundances are then set to zero and *cannot recover*. In large populations the dynamics and extinctions are driven deterministically by growth rate differences, so we neglect demographic noise. Starting from positive initial conditions, the strain abundances in the CR model generally reach a fixed point with some fraction of strains going extinct and others persisting at nonzero abundance. We denote by $\mathcal{C}$ the resulting *community of extant strains*: those with positive fixed-point abundances. For small enough S—less than some multiple of D—the fixed point is generally stable and cannot be invaded by any of the extinct strains. For large D, in the parameter ranges of interest, the stable uninvadable community is highly likely to be unique [40]. The breakdown of this uniqueness above a critical value of S in assembled communities (discussed for a related model in Ref. [22]), will not play a role in the evolving communities of interest here.

At the ecological fixed point of the CR dynamics, the abundances $\{n_i^*\}$ of the extant strains and the resources $\{\mathcal{R}_\alpha^*\}$ are determined (with $\omega = 0$) by

$$\mathcal{R}_\alpha^* = \frac{K_\alpha}{\sum_{j\in\mathcal{C}} \mathcal{F}_{j\alpha} n_j^*}, \tag{2a}$$

$$0 = \sum_{\alpha=1}^{D} \mathcal{G}_{i\alpha}\mathcal{R}_\alpha^* - d \quad \text{for } i \in \mathcal{C}, \tag{2b}$$

$$n_j^* = 0 \quad \text{for } j \notin \mathcal{C}, \tag{2c}$$

with the strains $j \notin \mathcal{C}$ extinct and unable to reinvade. Whether a strain can invade from low abundance is determined by the sign of its *invasion fitness*

$$\xi_i = \sum_{\alpha=1}^{D} \mathcal{G}_{i\alpha}\mathcal{R}_\alpha^* - d. \tag{3}$$

If ξ_i is negative, strain i cannot invade.

The total population size $N \equiv \sum_i n_i^*$ is of order $\mu_g \sum_\alpha K_\alpha/\mu_f$, which is roughly D for our choice of parameters. (We have kept the dependence on $\mathbf{K}$ explicit because it carries a physical meaning and characterizes the environment. In our simulations we have $K_\alpha = 1$ for all α: the effects of changing $\mathbf{K}$ are discussed briefly in Sec. VI.)

We define the *extant strain diversity*, $L = |\mathcal{C}|$, as the number of strains *left* in the community at the ecologically stable fixed point. In accordance with the competitive exclusion principle [28], the maximum L that can stably coexist at a

fixed point is D, as seen from the insufficiency of the fixed point conditions otherwise. For natural choices—including ours—of the ensemble of interactions, the typical number of strains that can stably coexist at a fixed point scales linearly with D with a coefficient less than 1. This linear relationship between the diversity of the community and the number of resource-mediating chemicals is expected to be generic across a range of consumer resource models with heterogeneous interactions, with the statistics of the interactions affecting the prefactor in the scaling [38]. In ecological language, the number of *potential niches* is D but not all of these can be filled.

3. Measures of diversity

In the ecological literature many different measures of diversity of communities have been introduced [51]. We use L, the total number of species in the community, which is sometimes called the *richness*. Other popular measures are the Shannon diversity: the weighted average of $\log(N/n_i)$ (throughout, logs are base e), and the *inverse Simpson index* (*inverse participation ratio* in physics terminology), given by $N^2/(\sum_i n_i^2)$. If there is a broad distribution of abundances in a community, these measures can be quite different. However, in our models, as we shall see, the distribution of abundances is not very broad so which measure we use does not affect our conclusions. For simplicity we refer to L as the "diversity."

B. Evolutionary processes

We now focus on slow evolution of stable communities in the CR model, assuming a clean separation of timescales between evolution and ecology. If the resource dynamics are fast enough, they can be integrated out to yield equations for the strain dynamics which reach the same fixed point irrespective of the resource timescale—we therefore assume fast resource dynamics to simplify our dynamical equations without affecting the evolutionary behavior. The ecological dynamics of the community run until a fixed point is reached. Then a single new strain is introduced, and the ecological dynamics again run until a fixed point is reached. The ecological timescale thus does not affect the evolutionary dynamics of primary interest.

1. Evolution by independent invaders

We first study the case of *independent invaders* into the community. This process is sometimes referred to as "bottom-up assembly," and will give us insight into evolutionary dynamics with large-effect mutations. We start with a single strain of consumer and a collection of D resources and repeatedly add an unrelated consumer strain drawn from the Gaussian phenotype ensemble. We define an *epoch* as the unit of time corresponding to a *successful* invasion followed by the relaxation of the ecological dynamics to a fixed point. Although the length of an epoch in ecological time varies between epochs, in our case of slow evolution the ecological timescale is not important.

With each new strain added, we find a stable uninvadable ecological fixed point of the community either by integrating the dynamical equations with a small extinction threshold, or by the more efficient iterative algorithm described in Appendix D. For sufficiently large D, we expect that, with high probability, this fixed point is the *unique* stable uninvadable fixed point for the collection of extant strains in each epoch—this intuition is based on analogy with the assembled community in which the uninvadable fixed point is unique as long S and L are of similar size [20], as is the case in our evolving system where the size of the extant community *plus* the new invader is analogous to S.

2. Components of invasion fitness

In order for a new strain, labeled I, to be able to invade, it must have positive invasion fitness ξ_I. Furthermore, if it successfully invades, its fixed-point abundance n_I^* will be approximately proportional to ξ_I, with the constant of proportionality determined by the properties of the community into which I invades, as can be seen by a perturbative analysis for large L. The dot product between the growth vector $\boldsymbol{\mathcal{G}}_I$ (the Ith row of $\mathcal{G}$) and the resource vector $\boldsymbol{\mathcal{R}}^*$ of the community determines the invasion fitness of strain I. Therefore, when we draw the phenotype of I—consisting of D-dimensional vectors $\boldsymbol{\mathcal{G}}_I$ and $\boldsymbol{\mathcal{F}}_I$—from their Gaussian ensemble, we condition on $\xi_I > 0$, which drastically speeds up numerics when the evolution pushes the extant distribution of interactions to have large mean, making invaders with $\xi_I > 0$ very rare [52].

The resource availability vector $\boldsymbol{\mathcal{R}}^*$ can be thought of as a modified version of the supply vector $\mathbf{K}$, where the availability of resource α is reduced from K_α by a factor of $\sum_j \mathcal{F}_{j\alpha} n_j^*$. Therefore, alignment between $\boldsymbol{\mathcal{G}}_I$ and $\mathbf{K}$ results in a contribution to the invasion fitness of strain I that is independent of the other strains in the community—this contribution to the invasion fitness is a form of *general fitness*, here defined as $U_i \equiv \sum_\alpha \mathcal{G}_{i\alpha} K_\alpha$, which captures the average performance of strain i against all of the D resources at the rate that they are supplied by the environment and thus scales with the magnitude of the resource supply vector. Therefore, ξ_I for an invader has a community-independent contribution proportional to the general fitness U_I as well as a community-dependent contribution from the ecological interactions. We refer to the community-dependent part of the invasion fitness as the *drive*, which will be defined more precisely for the models we introduce in Sec. IV.

It is useful for analysis to extend the definition of the invasion fitness to strains already in the community: for $i \in \mathcal{C}$ we define ξ_i as the growth rate strain i would have if it were removed from the community and then reintroduced at low abundance. This quantity arises naturally in the mean-field description of ecological models (Sec. V and Appendix I) and we sometimes refer to it as the *bias*, consistent with previous work [24]. The bias is therefore an invasion fitness appropriately defined for strains both inside and outside the community. Differences between biases of different strains come from both their general fitnesses and their drives, and such differences are responsible for the differential growth and extinction of strains.

As we show in Appendix C, in a randomly assembled community the *selective differences*—differences between general fitness terms U_i—are *first order* in the phenotypic differences, proportional to $\frac{1}{K_{\text{tot}}} \sum_\alpha \mathcal{G}_{i\alpha} K_\alpha - \langle \mathcal{G} \rangle$, while the variations in the drives are second order, proportional to

both $\frac{1}{K_{\rm tot}}\sum_\alpha \mathcal{G}_{i\alpha}K_\alpha - \langle\mathcal{G}\rangle$ and $\frac{1}{N}\sum_i \mathcal{F}_{i\alpha}n_i^* - \langle\mathcal{F}\rangle$, where we have defined $K_{\rm tot} \equiv \sum_\beta K_\beta$, and weighted average $\langle\mathcal{G}\rangle = \frac{1}{NK_{\rm tot}}\sum_{i,\alpha}\mathcal{G}_{i\alpha}n_i^*K_\alpha$ (with the equivalent definition for $\langle\mathcal{F}\rangle$). Therefore, the selective differences comprise the dominant contribution to the width of the bias distribution, larger by a factor of $\sqrt{D}$ than the contribution from the drives. For large D the selective differences determine both the ecological and early stage evolutionary dynamics and can, at least initially, strongly limit the diversity. Previous work [38] has forced the contributions to be comparable by scaling the mean and variance of $\mathcal{G}$ and $\mathcal{F}$ as $1/D$, while in fact it is their mean and *standard deviation* that should be the same order in a biologically motivated model. We later show that in a long-evolved community, the variations in the general fitnesses become much smaller, indeed comparable to the variations in the drives. However—importantly—this feature must be shown to emerge rather than assumed.

C. Eco-evolutionary dynamics

To explore the eco-evolutionary dynamics, we choose $D = 50$ resources with $\kappa = 0.4$ well away from the special value $\kappa = 1$ for which there is a Lyapunov function: we investigate that special case later. Starting from a single strain ($L = S = 1$), Fig. 1(a) shows the ecological dynamics of the strain abundances in Eq. (1) (with resources integrated out), punctuated by well-spaced evolutionary events when an invader comes in at the start of each new epoch. When the system has reached an ecological fixed point, a new strain is introduced, drawn randomly from the ensemble of independent strains conditioned on being able to invade (the conditioning speeds up numerics, since new strains which cannot invade will not affect the community). The new strain joins the community and the ecological dynamics find a new stable fixed point at which the abundances of the other strains are perturbed, with some strains possibly going extinct. Thereafter another strain is introduced. (One could weight the distribution of successfully invaders by a factor of their invasion fitness, as would arise from the establishment probability with demographic noise [35], but we ignore this effect, which would change some quantitative details of our analysis but would not affect qualitative conclusions.)

From now on, we focus on the evolutionary dynamics and ignore the ecological transients at the beginning of each epoch that drive some strains extinct and the others to a stable community. The changes caused by invading strains can be visualized over longer times by coarse-graining the time resolution from ecological to evolutionary epochs and tracking the fixed-point abundances of strains (calculated with the algorithm of Appendix D). Such strain abundance trajectories are shown near the beginning of the evolution in Fig. 1(b) and, after much evolution, in Fig. 1(c). We show only a few of the strains for ease of visualization.

Primary quantities of interest are the extant strain diversity L, the population mean general fitness $\langle U\rangle \equiv \frac{1}{N}\sum_i n_i U_i$ (with angular brackets denoting weighted averages over the community), and the probability that a randomly chosen strain can invade, denoted $p_{\rm inv}$. In addition to the community-level properties, we are interested in the statistical properties of strains including the distributions and temporal correlations

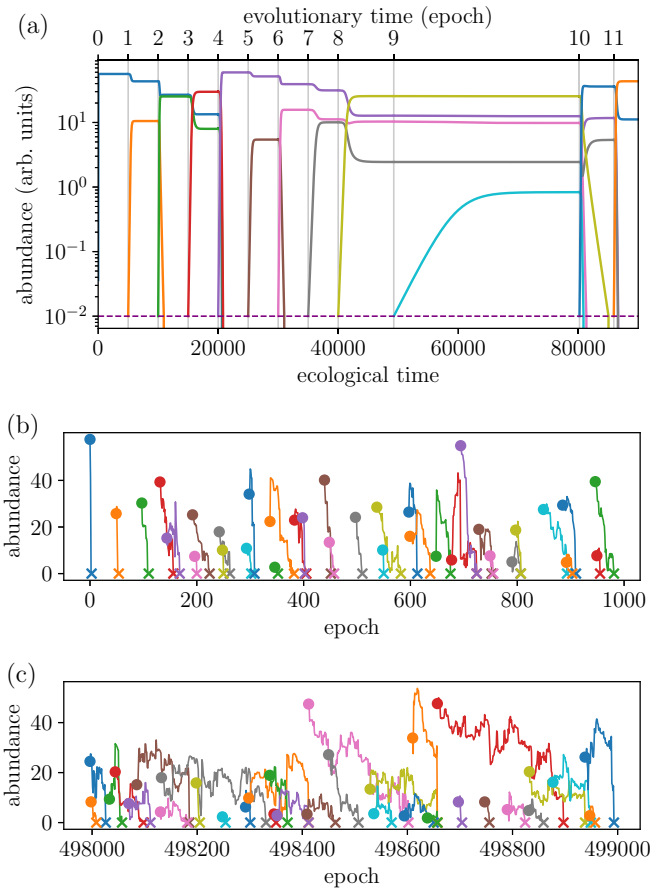


FIG. 1. Dynamics of strain abundances on ecological and evolutionary timescales in the consumer resource (CR) model for $\kappa = 0.4$ and $D = 50$ resources [Eq. (1)]. (a). Starting with a single strain, ecological dynamics reach a fixed point between each new strain invasion which begins a new epoch (vertical lines). Ecological time on lower axis; evolutionary time, in epochs, on upper axis. (b)–(c) Selected strain abundance trajectories on evolutionary timescales, with invasions marked by a circle and extinctions by a cross. (b) first 1000 epochs, (c) after a long period of evolution. Turnover of diversity is slower than during initial stages of evolution, but has converged to a steady average rate.

of their abundances n_i^*, their general fitnesses U_i and their evolutionary lifetimes which we call $\mathcal{T}_i$ (see Sec. V C for more details).

1. Asymmetric interaction ensemble

The evolution of the diversity for $\kappa = 0.4$ is shown in purple in Fig. 2(a). Starting from a single strain, L initially grows fast and then slows down, appearing to saturate at $L \approx 10$ at long times, for which we plot a moving average. As shown in Appendix M for a related model that we introduce shortly, the fluctuations in L are $O(\sqrt{L})$ resulting from rough independence of strains along with robustness of the community to perturbations (Sec. V B). For $\kappa = 0.4$, it thus appears that the evolution converges to a Red Queen steady state. Note that although for these parameters L remains rather small, we later show that it is expected to be of order $\kappa^2 D$, here only 8, in the large-D limit of primary interest. The modest $\kappa^2 D$ used

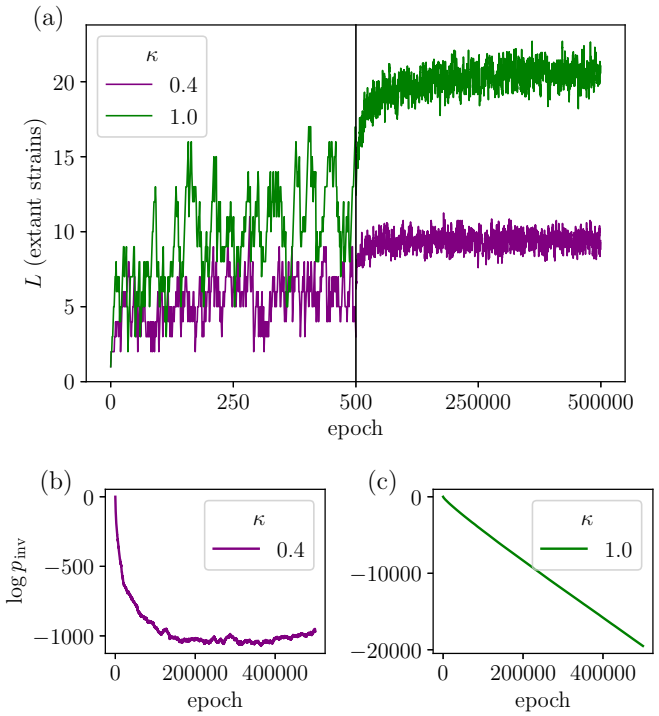


FIG. 2. The CR model reaches an evolutionary steady state for symmetry parameter $\kappa = 0.4$, but not for perfect symmetry, $\kappa = 1$. (a) Community diversity L as a function of evolutionary time, starting from a single strain with $D = 50$. First 500 epochs on stretched scale; thereafter as moving average over windows of 1000 epochs. L increases, then fluctuates around steady state for $\kappa = 0.4$, but creeps up for $\kappa = 1$. (b), (c) The invasion probability, p_{inv}, decreases initially, then reaches steady state for $\kappa = 0.4$, but continues to decrease for $\kappa = 1$. The absurdly small p_{inv} is an artifact of the Gaussian distribution of phenotypes (see Sec. V D).

here was needed for convergence to steady state in reasonable time.

Further evidence for the Red Queen state for $\kappa = 0.4$ is shown in Fig. 2(b), where the invasion probability of a new strain, p_{inv}, after initially decreasing rapidly, saturates at long times (albeit to an extremely small value). This behavior can be understood through the general fitnesses. Figures 3(a) and 3(b) show in purple the mean and width, respectively, of the extant U_i, distribution. Initially the mean increases, but then saturates at a value of order D with the strains all having similar U_i. The evolution therefore pushes far out into the tail of the general fitness distribution, making it very hard for other strains to invade. The width of the extant U_i distribution decreases initially and then saturates. We later show (Sec. V D and Fig. 6) that the width of the distribution of extant general fitnesses, scaled by the mean consumption rate of the resources—which is what enters in the growth rate of the strains (Appendix C)—shrinks during evolution until it is similar to the width of the drive distribution. At this point the intrinsic (general fitness) and community-derived (drive) parts of the bias contribute similarly to bias differences between strains. Once this occurs, the distribution of extant U_i stops pushing higher and we reach a Red Queen evolutionary steady state with a roughly steady rate of turnover.

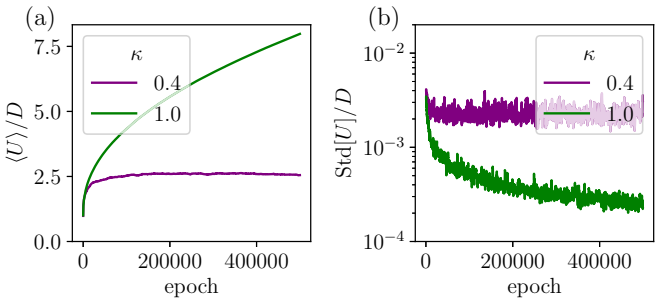


FIG. 3. General fitness of the community during evolution in the CR model. Mean (a), and standard deviation (b) (weighted by abundances) of extant distribution of generalized fitness, U_i, over evolutionary time for $\kappa = 0.4$ and 1 with $D = 50$. Distribution stabilizes for $\kappa = 0.4$. For $\kappa = 1$, mean continues increasing with standard deviation decreasing.

In an evolutionary steady state (when it occurs) we denote *averages over time* by overbars. Because of the turnover of strains, in steady state temporal averages should be equal to averages over the ensemble of strains.

2. *Turnover of strains*

In steady state, one new strain invades in each epoch and, on average, one extant strain goes extinct. This means that a strain remains in the community for an average of $\overline{L}$ epochs and the rate of turnover of strains is $1/\overline{L}$ per epoch. This is seen in Fig. 1 where at early times when L is small, strains do not survive long, while in the long-time steady state with larger $L \approx \overline{L}$ they survive longer on average even though there are still strains that go quickly extinct, as we discuss in Sec. V C.

The continual turnover of strains together with the saturation of p_{inv} mean that in the Red Queen phase, strains that successfully invade and later go extinct will have a substantial chance of being able to reinvade once the community has sufficiently turned over and lost memory of its state at the extinction time of the strain. Each strain keeps "memory" of its general fitness which, because it was able to invade earlier, must be similar to that of the other strains in the steady state.

While these simulations provide strong evidence for a Red Queen steady state for $\kappa = 0.4$, quantitatively they are problematic. It is apparent from Fig. 2(b) that, for the ensemble of strains we have used, the invasion probability of a new invader is unrealistically small in the steady state ($p_{\text{inv}} \approx 10^{-400}$). In Sec. IV A we introduce a family of models in which p_{inv} at steady state can be much larger, and in Sec. V D we analyze how the choice of the $\mathcal{G}$ and $\mathcal{F}$ ensembles affects the invasion probability in steady state and show that the very small p_{inv} observed above is an artifact of the Gaussian phenotype ensemble we have used for numerical convenience.

3. *Totally symmetric ensemble*

We now show that the perfectly symmetric case $\kappa = 1$, which possesses a Lyapunov function Λ (see Appendix A), behaves very differently than the $\kappa = 0.4$ ensemble for which there is no Lyapunov function. From the green trajectory of the diversity in Fig. 2(a), it is difficult, based only on the

dynamics of L, to know whether a steady state obtains by the end of the simulation. However, invading strains must increase Λ and there is no absolute upper bound of Λ in the ensemble of all possible invaders. Therefore, even if L saturates, there cannot be a steady state. The larger Λ gets, the more unlikely that adding new strains can increase it further—thus the probability of invasion, p_{inv}, must steadily decrease, as seen in Fig. 2(c). Concomitantly, the mean general fitness pushes further and further into the tail of its distribution as shown by the green trajectory in Fig. 3(a), while the standard deviation of the U_i continues to decrease [green curve in Fig. 3(b)] due to the steepening of the distribution around $\langle U \rangle$ (see Sec. V D).

Because of the Lyapunov function, we conclude that there is no Red Queen steady state for $\kappa = 1$. It is interesting to note that, in spite of the gradual increase of Λ, L stays far from its upper bound of $D = 50$. Indeed, L stays below its maximum value of $D/2$ for a random *assembled* community with $\kappa = 1$ [38]. However, it is not clear whether under evolution L might eventually reach or exceed this value because it is not a strict bound.

4. Evolution via mutations

So far we have considered independent invaders in the CR model. In Appendix B we study the evolutionary dynamics when invaders are correlated with a parent strain chosen with probability proportional to its abundance in the community. Indexing the parent by P and the mutant by M, we introduce a parameter $\rho \in [0, 1)$: the element-wise Pearson correlation coefficient between $\boldsymbol{\mathcal{G}}_P$ and $\boldsymbol{\mathcal{G}}_M$ and between $\boldsymbol{\mathcal{F}}_P$ and $\boldsymbol{\mathcal{F}}_M$. Given a parent, we draw mutants from this ensemble (while also preserving correlation κ between $\boldsymbol{\mathcal{G}}_M$ and $\boldsymbol{\mathcal{F}}_M$): therefore the correlation between U_P and U_M is also ρ. Our main finding is that for ρ sufficiently far from 1—i.e., mutations of substantial effect—the eco-evolutionary phenomenology is similar to the independent invader limit $\rho = 0$ (Fig. 8). This finding allows us to base qualitative conclusions for evolution by mutations on results from independent invaders.

IV. SIMPLIFIED MODELS

The saturation of the community-mean general fitness in the Red Queen eco-evolutionary steady state for $\kappa < 1$ is a manifestation of continual evolution in an ecological system. We seek simpler—but still robust—models that display such a Red Queen steady state, but which allow us to tune the relative magnitudes of ecological interactions and general fitnesses. The simplified model that we study (Sec. IV A) has a generalized Lotka-Volterra form that comes from integrating out the resource dynamics and linearizing the resource dependence on the consumer abundances, yielding a maximum rank D for the effective interaction matrix.

We have also studied full-rank Lotka-Volterra (LV) models with more standard random interaction matrices of heterogeneous pairwise interactions possessing a tunable degree of symmetry, and diagonal "niche" interactions which enable substantial diversity to coexist (Appendix F). The fully asymmetric version of this model with no correlations between pairwise interactions emerges in a particular high-D limit of the linearized CR model (Sec. V C) and is the basis for the exact analyses we carry out. In Sec. VI we describe a qualitatively different phase of eco-evolutionary dynamics that appears in the LV model as the symmetry of its interaction matrix increases: in this oligarch phase an $O(1)$ fraction of the community's abundance condenses into an $O(1)$ number of strains, and the speed of evolution is ever-decreasing.

A. Linearized resource model

In the Red Queen steady state of the CR model, as seen in Appendix C, the mean of $\sum_j \mathcal{F}_{j\alpha} n_j^*$, the total consumption rate of resource α, is large compared with its variations, suggesting an expansion in these variations. In addition, the total population of the ecosystem $N = \sum_j n_j^*$ is roughly constant on evolutionary timescales, due to the large mean of $\mathcal{G}$ and $\mathcal{F}$ relative to their variations. Therefore, it is helpful to separate out the mean part of these matrices from their variations and introduce zero-mean residual interaction matrices G and F. We also work in terms of fractional abundances $\nu_i = n_i/N$ rather than absolute abundances and integrate out the resource dynamics since these do not affect the ecological fixed points.

Expanding in the *differences* between resources and between strains yields general fitnesses at first order and interactions at second order (Appendix C). The simplified model that we introduce, and refer to hereafter as the *linearized resource model* or simply the linearized model, is

$$\frac{d\nu_i}{dt} = \nu_i \left(s_i + \sum_{\alpha=1}^{D} G_{i\alpha} R_\alpha - \Upsilon(t) \right), \tag{4a}$$

$$R_\alpha = -\sum_{j=1}^{S} F_{j\alpha} \nu_j. \tag{4b}$$

The term $\Upsilon(t)$ is a Lagrange multiplier, equal to $\sum_i \nu_i (s_i + \sum_\alpha G_{i\alpha} R_\alpha)$, which incorporates the average effects of growth, death, and resource depletion, and keeps the population constant with $\sum_i \nu_i = 1$. Here, the per capita growth rate of each strain is a linear function of the abundances of all the other strains, with a quadratic part from Υ that is the same across all the strains and includes the average general fitness. Differences in general fitness are parametrized by s_i, the *selective differences* between strains, which play the role of the U_i in the CR model. Note that the "resource vector" $\mathbf{R}$ can have components which are negative; $\mathbf{R}$ then measures *deviations* of resource availability away from some positive baseline.

One could parametrize the s_i as $s_i = \sum_\alpha G_{i\alpha} k_\alpha$ for some fixed "supply vector" $\mathbf{k}$ in which case we could write $R_\alpha = k_\alpha - \sum_j F_{j\alpha} \nu_j$ and eliminate the separate s terms, but we focus on the case where the s_i and corresponding row of G are independently drawn, which does not change the behavior much in the regimes we study.

The structure of Eq. (4) is that of a generalized Lotka-Volterra model or system of replicator equations (discussed further in Appendix F), where the interaction matrix between S strains, $\mathcal{V}$, has a structure inherited from the G and F matrices via $\mathcal{V} = -GF^\intercal$, where G and F are $S \times D$ matrices so that $\mathcal{V}$ is rank D. We take each element of G and F to be Gaussian distributed with mean zero and unit variance, with no correlations within G or F. Element-wise correlations between G and F are again parametrized by $\kappa \in [0, 1]$ via $\mathrm{E}[G_{i\alpha} F_{j\beta}] = \kappa \delta_{ij} \delta_{\alpha\beta}$.

The off-diagonal entries of $\mathcal{V}$ have mean zero and variance D while the diagonal entries have mean $-\kappa D$ and variance $(1+\kappa^2)D$. The parametrically larger (for large D) diagonal entries have negative-mean larger than the standard deviation of the off-diagonal elements by a factor of $Q \equiv \kappa\sqrt{D}$: this endows the interactions with sufficient *niche structure* to stabilize ecological communities. (Note that $\kappa < 0$ would result in unstable behavior with the diagonal entries of $\mathcal{V}$ having positive mean.) As in the CR model, the number of resources D sets the dimensionality of phenotype space, with the phenotype of strain i comprised of s_i along with the ith rows of G and F. The maximum number of coexisting strains at a stable fixed point of the linearized model is equal to D and will generally be proportional to D with a prefactor less than 1.

In the special case $\kappa = 1$, Υ is closely related to a convex Lyapunov function for the dynamics of Eq. (4): the Lyapunov function is $\mathcal{L} \equiv \sum_i \nu_i(2s_i + \sum_\alpha G_{i\alpha}R_\alpha) = \Upsilon + \langle s\rangle$ where $\langle s\rangle \equiv \sum_i \nu_i s_i$. If all $s_i = 0$, we have $\Upsilon = \mathcal{L} = -\sum_\alpha(\sum_i G_{i\alpha}\nu_i)^2$ which is nonpositive. For nonzero s_i, the Lyapunov function gets an extra piece $2\langle s\rangle$ and can be of either sign, although it remains a convex function of the $\{\nu_i\}$.

We take the general fitnesses s_i to be drawn from a distribution $p(s)$, which for most of this paper is Gaussian with mean 0 and variance Σ^2. In Sec. V D we show that the value of Σ determines the early evolutionary behavior and the value of p_{inv} at steady state, with the shape of the large-s tail of $p(s)$ playing a controlling role. As the community evolves and the distribution of extant s_i pushes into the tail of $p(s)$, the variance of the extant s_i decreases. Therefore, tuning Σ to be small *a priori* mimics the effects of prior evolution and allows us to access long-time evolutionary behavior without a long initial transient.

As discussed previously in Sec. III B, for the CR model in Eq. (1), the width of the U_i distribution is a factor of $\sqrt{D}$ larger than the width of the drive distribution. In the linearized model, this corresponds to $\Sigma = O(\sqrt{D})$, since the distribution of the drives in a diverse community has $O(\sqrt{D/L}) = O(1)$ width. However, in the Red Queen steady state, as we shall see, the first and second order (in $1/\sqrt{D}$) components of the growth rate of each strain become comparable to one another. Thus we mostly consider $\Sigma = O(1)$, and sometimes set $\Sigma = 0$, which shortens the transient period of evolution before the steady state is reached—and enables analysis—but does not change our conclusions. In Sec. V D we also discuss non-Gaussian distributions of the s_i, which affect quantitative properties but not the basic phenomenology.

B. Evolution in the linearized resource model

The linearized model captures the phenomenology of the CR model—a Red Queen state with continual turnover but roughly constant mean general fitness and p_{inv}—in a simple setting where the effects of the general fitnesses can be separated out from the interactions.

As in the CR model, whether a new strain I successfully invades the extant community or not is determined by its invasion fitness or bias, now written as $\xi_I = s_I + \zeta_I - \Upsilon$ with $\zeta_I = \sum_j \mathcal{V}_{Ij}\nu_j$. Splitting the bias up into the general fitness piece s_I, the Lagrange multiplier Υ and the drive ζ_I is useful since it separates the intrinsic and extrinsic contributions to the bias. Strains with positive bias can invade the community and reach abundance proportional to their bias, whereas strains with negative bias go extinct. As in the CR model, we sample properties of the invader I conditional upon $\xi_I > 0$ [52]. A strain's extinction occurs when its abundance goes to 0 in the ecological dynamics or, equivalently, when its bias goes negative.

In general, the evolutionary dynamics of the linearized model [Eq. (4)] are parametrized by κ, Σ, the mutant-parent correlation ρ, and D which sets the overall scale of diversity. The initial pool size S only plays a role in the transient behavior: the evolution converges to the same state independent of the initial diversity [see Fig. 4(a)].

Simulations of the evolutionary dynamics of the linearized model across ranges of κ, ρ, and Σ (Appendix E) show that the Red Queen phase is very robust, obtaining across a range of parameters for $\kappa < 1$, with parameter choices affecting the length of the pre-steady-state transient and the quantitative properties of the steady state. Thus the Red Queen state can be qualitatively understood by studying just a few values of the relevant parameters. The value of D does not matter much as long as it is large, and the value of ρ does not matter much as long as it is not too close to $\rho = 1$. We also see that the steady-state diversity is independent of Σ, but the length of the transient increases with Σ and with κ.

V. ANALYSIS

To understand the phenomenology of the Red Queen phase described above, it is useful to first discuss the *top-down assembled* community of the linearized model in which all strains are brought together at once with dynamics determined by their ecological interactions. An understanding of the ecological phase diagram in this setting is informative for analyzing the Red Queen phase and gives us an idea of how $\overline{\Upsilon}$ and $\overline{\log p_{\text{inv}}}$ scale with various parameters in the steady state.

To analyze the eco-evolutionary dynamics, in Sec. V C we introduce a particular simple limit of the linearized model which reduces to a full-rank totally *asymmetric* LV model with no correlations in its interaction matrix. With the further simplifications of independent invaders and $\Sigma = 0$—yielding the simplest ensemble of the LV model—we can exactly solve for the steady state of the evolutionary dynamics in the limit of large diversity. Key quantities are the distribution of strain abundances in an evolved community—strikingly different than a top-down assembled community—and temporal correlation functions of strain abundances in evolutionary time.

In Sec. V D we then add in selective differences, presenting general scaling arguments for the dependence of $\overline{L}$, $\overline{\Upsilon}$, and $\overline{\log p_{\text{inv}}}$ on Σ, which build on the intuition gleaned from the assembled communities. The presence of selective differences changes the temporal decay of strain abundance correlation functions, and we discuss subtleties associated with long-lived strains with anomalously large s_i. We finally generalize to different distributions, $p(s)$, and show how the primary pathologies of the Gaussian distribution can be alleviated.

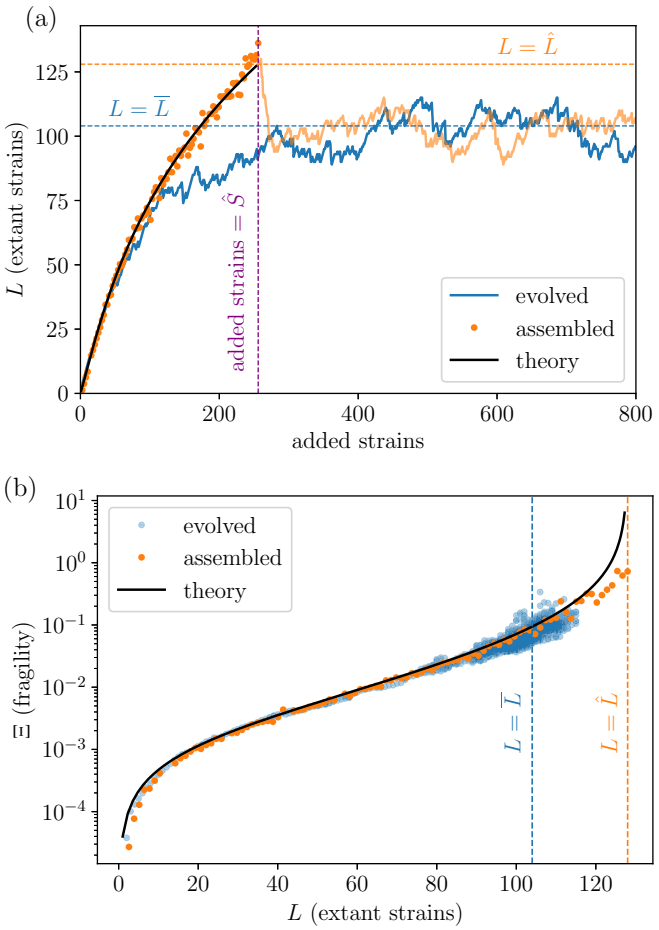


FIG. 4. Comparison of evolved and assembled communities showing lower diversity in evolved communities due to perturbations from invading strains. Linearized model with $\kappa = 0.8$, $D = 200$, $\rho = 0$, and $\Sigma = 0$. (a) Number of extant strains, L, as a function of number of strains added. Blue line: evolving community with strains added one by one and permanent extinctions after addition, showing convergence to Red Queen steady state with average diversity $\overline{L}$. Yellow points: strains assembled all at once (or added one-by-one with no extinctions); simulations match the black line from the assembled cavity calculation (Appendix I). Light-yellow line: evolutionary dynamics, with extinctions, starting from saturated assembled community, which initially has $L > \overline{L}$ but converges to the same Red Queen state. (b) Fragility Ξ, a community-wide measure of susceptibility to perturbations, versus L [Eq. (I12)]. Fragility diverges at transition between stable and unstable phases (vertical dashed line), which is not reached in evolved communities (blue points).

A. Assembled communities

For a top-down assembled community in the linearized model, all strains are initialized at positive abundance and are governed by the deterministic dynamics of Eq. (4) without an extinction cutoff. The cavity method [53] can be used to analyze both the statics and dynamics of the assembled community with S initial strains and D resources, in the joint limit of large S and D with fixed S/D. The spirit of the cavity method—in the static case—is to consider a large community to which one adds, separately, one new strain and one new resource: these each act as small perturbations to the fixed point of the community. The abundance of the new strain is determined by its invasion fitness, modified by the community-wide response to its introduction, and the new resource is treated similarly. One then enforces self-consistency conditions that the statistics of the strains and resources in the community are the *same* as the statistics of the newly added strain and resource. These conditions rely on the community being assembled with independent interactions, since then the independently drawn invading strain and resource are from the same ensemble as the rest of the community. By solving the self-consistency equations one can extract abundance distributions of ν_i and R_α, and solve (in the linearized model) for L/D and Υ as a function of S/D, κ, and Σ. Similar results can be obtained for the CR model [38] and—more simply—for Lotka Volterra models [21].

The cavity method also allows one to derive self-consistent equations for the ecological *dynamics* [53], in addition to their fixed point. We use a variant of this method to understand the *eco-evolutionary* dynamics in Sec. V C.

Details of the static cavity analysis for top-down assembled communities with Gaussian $p(s)$ are in Appendix I. This analysis shows that generically the linearized model exhibits two phases as a function of the ratio S/D. For $\Sigma > 0$, when S/D is small enough, there is a unique stable globally attracting fixed point, uninvadable by the strains that are going extinct: this fixed point is described by the cavity solution and the community diversity $L < \kappa^2 D$. The abundances of the strains in the stable community are distributed as a truncated Gaussian so that there is a constant density of abundances near zero. The strains destined for extinction have negative invasion fitnesses with a distribution that also has a constant density near zero. Thus the assembled community has many close-to-marginal strains.

When the number of the strains assembled, S, becomes larger than a critical value $\hat{S}(D, \kappa, \Sigma)$, at which $L = \hat{L} = \kappa^2 D$, the cavity solution becomes unstable and chaotic dynamics or multistability ensues [22]. We use a hat to denote the values of quantities at this phase transition and refer to the community at this transition as *saturated*. The number of initial strains needed to achieve this saturated state, $\hat{S}$, diverges as $\kappa \to 1$ as long as $\Sigma > 0$. Therefore, for $\kappa = 1$ and $\Sigma > 0$ the cavity solution is always stable. [When $\Sigma = 0$ the picture for $\kappa < 1$ remains the same, while at $\kappa = 1$ there is a transition (for $S > 2D$) to a nongeneric "shielded" phase [37] in which $L/D = 1$ (the maximum allowed by competitive exclusion) and this state is highly marginal.]

With an extinction threshold, even for $S > \hat{S}$ the ecological dynamics will generally reach a stable fixed point with $L \lesssim \kappa^2 D$. However, the resulting community will generically be unstable to re-invasion by some of the strains that had gone extinct. Furthermore, in this supersaturated regime there are multiple possible stable communities, almost all of them invadable by some of the extinct strains. Which community occurs will thus depend on the initial conditions. The question of whether or not stable *uninvadable* communities are likely to exist for large D is unresolved except for the special $\kappa = 1$ case where there is an absolute maximum of the convex Lyapunov function and there is always a unique globally attracting stable community. It is believed that, even if stable

uninvadable fixed points exist for $\kappa < 1$, they may attract only an exponentially small fraction of initial conditions [18,54].

We can glean insights about the properties of an evolved community with $\rho = 0$ (independent invaders) from an analysis of the near-saturated assembled community, since after long period of evolution, a large number of strains—many more than $\hat{S}$—will have had a chance to invade the evolved community. Naively, one might think that the end result will be similar to assembling the whole community at once. However, as we shall see, the evolved community never becomes saturated in the way that an assembled community can be.

We would like to understand the behavior for $\Sigma \gg 1$ which is what occurs "naturally" (e.g., from the CR model): we thus focus for now on this limit. In the linearized model the saturated assembled community has $\hat{L} = \kappa^2 D$ for all Σ and, for $\Sigma \gg 1$,

$$\hat{\Upsilon} \approx \frac{\Sigma^2 \kappa}{1-\kappa^2}, \tag{5a}$$

$$\hat{S} \sim D \exp\left[\frac{\kappa^2 \Sigma^2}{2(1-\kappa^2)^2}\right]. \tag{5b}$$

Therefore, in the saturated assembled community, the invasion probability of a single independent invader is the fraction of surviving strains, which is $\hat{L}/\hat{S}$.

We conjecture that the *forms* of the scalings of $\hat{\Upsilon}$ and $\log \hat{S}$ with Σ in the assembled community also describe the time-averaged quantities in the Red Queen steady state, i.e. that $\overline{\Upsilon} \sim \hat{\Upsilon}$ and $\overline{\log p_{\rm inv}} \sim \log(\hat{L}/\hat{S})$, albeit with unknown coefficients. (Note that the *typical* steady state $p_{\rm inv}$ is better described by $\overline{\log p_{\rm inv}}$ than by $\log \overline{p_{\rm inv}}$, since $p_{\rm inv}$ can fluctuate by large amounts, particularly when Σ is large, as discussed in Sec. V D.) The scalings in Eq. (5), along with the guess that $\overline{L} \sim \hat{L}$, motivate the approximate theory curves in Appendix E: these are consistent with numerical data from the Red Queen phase. In Sec. V D we give an argument that the conjectured scalings of $\overline{\Upsilon}$ and $\overline{\log p_{\rm inv}}$ with Σ should indeed be expected on evolutionary grounds.

B. Fragility of assembled and evolved communities

In the eco-evolutionary Red Queen phase, ecological communities remain stable, but are different from top-down assembled communities since their interaction statistics are conditioned on evolution and, crucially, they would be *unstable to invasion* by some fraction of previously extant strains that have gone extinct. An indication of this distinction, shown in Fig. 4(a), is the diversity of the Red Queen phase: $\overline{L}$ is smaller than the maximum L of a saturated assembled community. When a community is assembled one strain at a time with *permanent* extinctions, it is not possible to reach as high diversity as reached by top-down assembly. What is responsible for this gap between $\overline{L}$ in the Red Queen phase and $\hat{L}$ from the saturated assembled community? The answer is associated with the relative fragility of the two states.

1. Fragility of assembled communities

The transition between phases in the assembled community is marked by the onset of an instability in the cavity solution. This instability entails the divergence of a community-level property that we call the *fragility* and denote by Ξ. The fragility quantifies the variance of the changes in strain abundances in response to a random perturbation to the growth rates of all the strains. Concretely, if we add a small independently random term h_i to the growth rate of each strain i, with $\mathrm{E}[h_i] = 0$ and variance $\mathrm{E}[h_i^2]$ (loosely like applying random fields in magnetic systems), and we denote by $\delta\nu_i$ the change in the abundance of strain i in response to this global perturbation, the fragility is defined by $\sum_i \mathrm{E}[\delta\nu_i^2] = \Xi\ \mathrm{E}[h_i^2]$. The fragility can also be written as $\Xi = \mathrm{Tr}[\chi^\intercal \chi]$ where the $\chi_{ij} = \frac{d\nu_i}{d\xi_j}$ are elements of the susceptibility matrix χ of the community.

Deep in the stable phase the fragility is low, as small perturbations to strains' growth rates do not cause abundances to reshuffle by much. However, as the saturation transition of an assembled community is approached where L nears $\hat{L} = \kappa^2 D$, the fragility diverges, indicating a breakdown in the stability of the cavity solution. In Fig. 4(b) we plot the fragility, computed directly from the interaction matrix, for both an evolving and a top-down assembled community across a range of L. In both cases, Ξ increases with L, although in the evolving community L plateaus below $\hat{L}$ and the fragility does not get close to diverging.

2. Fragility as response to invaders

The role of the fragility in distinguishing the assembled and evolved communities can be understood from its interpretation as a response to evolutionary perturbations. Each new invader adds a small random part to the biases of the other strains. Invader I, if successful, reaches a stable abundance ν_I, and contributes $\mathcal{V}_{iI}\nu_I$ to the bias of each other strain, i. The variance of this random perturbation is $\mathrm{Var}[\mathcal{V}_{iI}\nu_I] = \nu_I^2 D$ since the off-diagonal elements of $\mathcal{V}$ have typical magnitude $\sqrt{D}$. The mean-square effects of this random perturbation are multiplied by a factor of $\Xi/L \sim 1/L^2$ to yield $\mathrm{Var}[\delta\nu_i]$ of order $1/L^3$ and therefore $\delta\nu_i \sim L^{-3/2}$. For a typical strain with $\nu_i \sim 1/L$, this perturbation will be small. However, in a saturated top-down assembled community, which has a truncated Gaussian distribution of abundance with constant density as $\nu \to 0$, the lowest-abundance strains will have ν_i smaller than typical by a factor of $1/L$, and $O(\sqrt{L})$ of these will have abundance less than the magnitude of the perturbation. This implies that the invasion will drive the bias of $O(\sqrt{L})$ of the lowest-abundance strains negative, causing their extinction. As long as Ξ is not large, this number of extinctions will be small. But as $S \nearrow \hat{S}$, Ξ diverges and the assembled community will be shaken a lot by a single additional strain.

3. Fragility of evolved communities

The behavior of an *evolved* community—with permanent extinctions—in response to a small random perturbation is quite different. As stated above, for a near-saturated assembled community, a single invader will drive a fraction of the strains with $\nu \leqslant O(L^{-3/2})$ extinct. As this is $O(\sqrt{L})$ strains, L will initially decrease [light-yellow trajectory in Fig. 4(a)]. Under the perturbations from further invasions, the density of low-abundance strains will be hollowed out by this process, which is like diffusion with an absorbing boundary condition at zero: This results in a density of abundances that goes

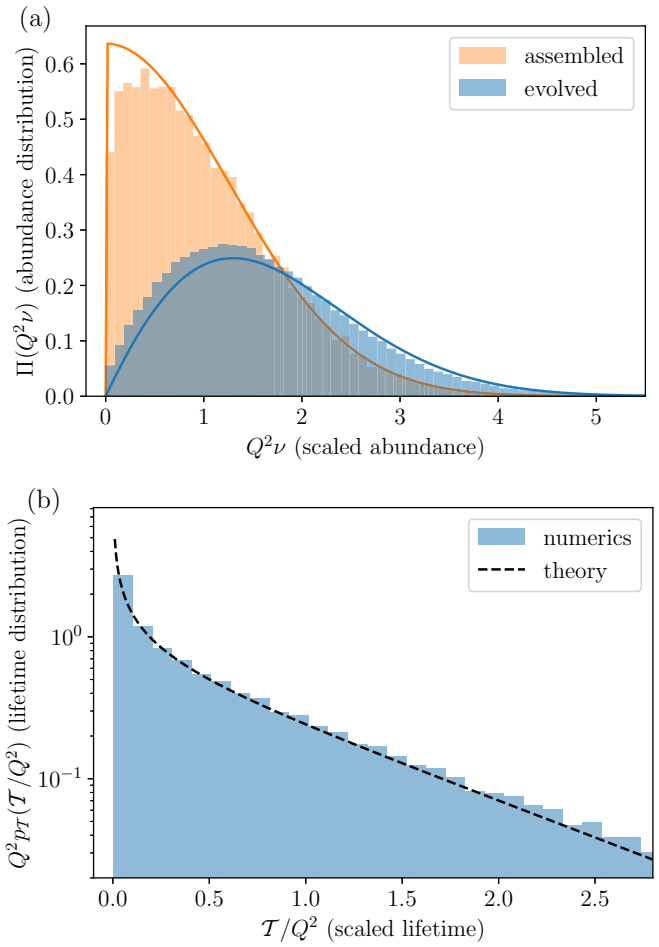


FIG. 5. Distributions of abundances (a) and lifetimes (b) in the Red Queen steady-state of the fully asymmetric LV model with $\Sigma = 0$ and $Q = 20$ (here, Q^2 corresponds to $\kappa^2 D$ in the linearized CR model). (a) Number density $\Pi(z)$ of the scaled abundance $z = Q^2\nu$: assembled communities in yellow, and evolved Red Queen steady state in blue with solid lines predictions from, respectively, assembled cavity calculation (Appendix I) and evolutionary cavity calculation [Eq. (K13), with normalization $\int dz \Pi(z) = \bar{L}/Q^2$]. (b) Distribution of strain lifetimes $\mathcal{T}$, normalized by basic timescale Q^2, agrees with prediction of Eq. (11) (solid line).

linearly to zero as $\nu \to 0$ [see Fig. 5(a)]. The diffusion coefficient of the abundances in evolutionary time is proportional to the fragility of the community, so if Ξ is large, then the perturbations are amplified and more extinctions result, on average, from each successful invader. As the diversity decreases, so does its fragility until a steady state is reached with on average one extinction driven by each invasion.

On the other hand, starting from a single strain, as the diversity grows we expect that (as in the top-down assembled communities), the fragility also grows. When it is large enough that on average one extinction is driven by each invasion, the diversity stabilizes, as shown in Fig. 4. This yields the *same* steady state reached upon decreasing from the saturated assembled community. In this steady state, the fragility is not too large suggesting—although not implying because of the complexities on conditioning on the evolutionary history—that the Red Queen state will be less diverse than the maximal assembled diversity: Indeed, $\bar{L} < \hat{L} = \kappa^2 D$, as seen in Fig. 4(a). (We note that the ratio $\bar{L}/\kappa^2 D$ depends on the value of κ as seen in Fig. 10(b); this is due to the dependence of Ξ on κ.)

The fact that a top-down assembled community (if close to saturated) can be larger than the evolved community means that there are naturally findable stable ecological communities that are unstable to evolution, in which L decreases as the evolving community converges to the Red Queen phase. On the other hand, top-down assembled communities can also be as small as a single strain, in which case their diversity increases under evolution. Therefore, even in simple models, the question of whether evolution increases or decreases diversity from an assembled initial condition has no general answer [29,31].

In the special case of perfect symmetry ($\kappa = 1$), L approaches $\hat{L}$ at long times despite the lack of a Red Queen steady state. We do not have an explanation for this apparent saturation of the diversity through gradual assembly.

In recent work, de Pirey and Bunin [50] have observed and studied a similar phenomenon of depletion of low-abundance strains (in their case only partial depletion) due to continual perturbations to the community from invading strains. These authors looked at a purely *ecological* model with small migration from a mainland that prevents total extinction of any strain. Species rising up from the brink of extinction play the role of invaders in our evolutionary model. The similarities between this ecological model and our work are discussed in Appendix P.

C. Exact solution for evolutionary dynamics in a simple limit

Having understood some features of the eco-evolutionary phenomenology by comparison with the saturated assembled community, we now proceed to directly analyze the evolutionary dynamics for a particularly simple choice of parameters.

With $\kappa > 0$, there are correlations κ^2 between the cross-diagonal elements of the interaction matrix $\mathcal{V}$. These correlations induce complex memory into the evolutionary dynamics since the effect of an invader I on a strain j in the community via $\mathcal{V}_{jI}$ feeds back via $\mathcal{V}_{Ij}$ onto I. This feedback on I is coherent across the extant strains with magnitude proportional to κ^2. Thus the properties of a strain in the community depend on its effects on the other strains over all the epochs since its invasion. This fact makes the evolutionary dynamics difficult to analyze (like the ecological dynamics in Ref. [24]), as the form of this memory kernel needs to be self-consistently determined along with the dynamics of the strains abundances in evolutionary time. We set this analysis up in Appendix L. The resulting self-consistency equations are more complicated and we have not attempted to solve them, although in principle this could be done numerically.

1. Fully asymmetric Lotka-Volterra model

To make progress, we analyze a simple limit of the linearized model: taking $D \to \infty$ and $\kappa \to 0$ while keeping the product $Q \equiv \kappa\sqrt{D}$ finite and rescaling the interaction matrix $\mathcal{V}$ by a factor of $\sqrt{D}$ leading to a Lotka-Volterra model with independent Gaussian-distributed interactions. All elements of the resulting scaled interaction matrix $V \equiv \mathcal{V}/\sqrt{D}$ have

unit variance. But the diagonals of V have mean $-Q$ and the off-diagonals have zero mean. The cross-diagonal correlations of V vanish in this limit, so we refer to this model as the *fully asymmetric Lotka-Volterra* (LV) model. Q^2 roughly plays the role of D in the linearized model, with the maximum diversity of the community scaling with Q^2 (although, in contrast to CR models, there is no strict upper bound). The lack of correlations within V makes the feedback effects smaller by a factor of $1/\sqrt{L}$ (and of random sign) and hence negligible—the analysis thus becomes tractable.

We later consider a broader class of random LV models, parametrized by a symmetry parameter $\gamma \in [-1, 1]$ which quantifies the cross-diagonal correlations in the interaction matrix: here it is equal to zero but some effects of nonzero γ are described in Sec. VI with more details in Appendix H.

2. Dynamic cavity method for fully asymmetric Lotka-Volterra model

Our analysis of the Red Queen phase proceeds by means of a dynamical cavity method as used in systems with quenched randomness [53]. The idea is to consider the evolutionary-time dynamics of a focal "cavity strain" labeled 0 that could potentially enter the community, persist for some number of epochs and later go extinct. The abundance of the cavity strain while it is in the community is determined by its bias, as long as this is positive. The strain's feedback on itself averages to zero so that its abundance is simply $\nu_0 = \xi_0/Q$ in the limit of large L (Appendix I). However, the bias, $\xi_0(T) = s_0 + \zeta_0(T) - \Upsilon(T)$, changes in evolutionary time T. If the strain attempts to invade at epoch $\mathbb{T}_0$, then its abundance in the community at all later times is

$$\nu_0(T) = \frac{\xi_0(T)}{Q} \mathbb{1}[\xi_0(T') > 0 \,\forall\, T' \in (\mathbb{T}_0, T)], \tag{6}$$

where $\mathbb{1}[E]$ is the indicator function which is unity when its argument is true and zero otherwise. The cavity strain's abundance is only nonzero as long as its bias remains positive—it goes irrevocably extinct once its bias first goes negative. Therefore, the statistics of the cavity strain abundance are determined by the statistics of its bias together with the permanent extinction condition.

To carry out the analysis, one needs to make an *Ansatz* for the statistics of the cavity bias, and then calculate the statistics of the cavity abundance as determined by Eq. (6). One then enforces self-consistency, which relates the dynamics of strain abundances in the community to the original statistics of the cavity bias. This is possible since the cavity drive is simply $\zeta_0(T) = \sum_{j\in\mathcal{C}(T)} V_{0j}\nu_j(T)$, where $\mathcal{C}(T)$ is the extant community at time T. Therefore, $\zeta_0(T)$ has Gaussian correlations for large L. Since the V_{0j} are $O(1)$, and are independent and unconditioned on properties of the community, the correlation function of the cavity drive is

$$C(\tau) \equiv \overline{\zeta_0(T)\zeta_0(T+\tau)} = \sum_j \overline{\nu_j(T)\nu_j(T+\tau)}, \tag{7}$$

where the sum is over *all* strains that have invaded before time T—most of which will be extinct at T or $T+\tau$ and therefore not contribute to the correlation function. For large $\overline{L}$, the average over T, denoted by the overbar, is not needed: we expect $C(\tau)$ to be "self-averaging" and depend only on the time-difference τ (up to subdominant corrections). A particular role is played by the equal-time correlations: $\sigma_\zeta^2 \equiv \mathrm{Var}[\zeta_I] = C(0) = \sum_i \nu_i^2$. (Note that, in the linearized model, σ_ζ carries an additional factor of $\sqrt{D}$ from the magnitude of the elements of $\mathcal{V}$.)

The self-consistency condition at the heart of the cavity method is that the autocorrelation of the cavity strain, averaged over trajectories $\{\zeta_0(T)\}$ and invasion time $\mathbb{T}_0$, is the same as the autocorrelation function of strains i in the community conditioned on having the same general fitness $s_i = s_0$. The correlation functions of strains in the community, weighted-averaged over the s distribution, determine $C(\tau)$ via Eq. (7). The Lagrange multiplier $\overline{\Upsilon}$ has to be adjusted to fix $\sum_i \nu_i = 1$ (just as Υ is determined for the static ecological cavity analysis of Appendix I).

We now specialize to the simplest case where all s_i are zero. Simulations of the fully asymmetric LV model with $\Sigma = 0$ [Fig. 13(b)] suggest that $\overline{\Upsilon} = 0$ in the steady state and we make the *Ansatz* that this is exact. ($\overline{\Upsilon} = 0$ implies a marginal ecological perturbation associated with changing the total population, but since our total population is fixed this does not imply marginality in the dynamics.) Hoping for some luck, we then make the simplest possible *Ansatz* for the statistics of $\zeta_0(T)$: namely, that it has an exponentially decaying correlation function $C(\tau) = \sqrt{J/b}\, e^{-b\tau}$. This *Ansatz* for the correlation function can be translated into a stochastic process that determines ζ_0:

$$\frac{d\zeta_0}{dT} = -b\zeta_0 + \sqrt{2J}\eta_0(T), \tag{8}$$

with T the time measured in epochs and $\eta_0(T)$ white noise satisfying $\langle \eta_0(T)\eta_0(T')\rangle = \delta(T-T')$ (angular brackets here denote averages over the stochasticity caused by invaders). The white noise represents the changes in the community from adding independent invaders which cause the diffusive changes in abundances discussed in Sec. V B. The deterministic term is, roughly, regression to the mean ζ_0 of zero, which can be seen to arise from the correlation between ζ_0 and $\delta\zeta_0$ due to the same elements of $\mathcal{V}$ appearing in both quantities [45]. The quantities b and J are (for now) phenomenological parameters that must be fixed later by self-consistency. The advantage of having a Markovian stochastic differential equation for ζ_0, along with $\overline{\Upsilon} = 0$, is that the method of images can be used to find the statistics of ν_0 with an absorbing boundary condition at zero, conditioned on its invasion at time $\mathbb{T}_0$.

To calculate the right-hand side of the self-consistency condition, Eq. (7), we need to calculate the two-time correlation function of ν_0. Since (with all $s_i = 0$) the strains are statistically equivalent, we can calculate $\langle \nu_0(T)\nu_0(T+\tau)\rangle$ for the cavity strain entering the community at time $\mathbb{T}_0$ with the initial value of $\zeta_0(\mathbb{T}_0)$, averaged over the stochasticity in Eq. (8). We then average this correlation function over $\mathbb{T}_0$ and $\nu_0(\mathbb{T}_0)$. We know that $\nu_0(T)$ is determined by Eq. (6) with $\xi_0(T) = \zeta_0(T)$ since $s = \overline{\Upsilon} = 0$.

From our *Ansatz*, the initial $\nu_0(\mathbb{T}_0)$ is distributed as a Gaussian random variable with standard deviation $\sigma_\zeta/Q = \frac{1}{Q}\sqrt{J/b}$, conditioned to be positive (as epochs are measured per successful invasion). To calculate $C(\tau)$ from the strains in

the community, we integrate $\langle \nu_0(T)\nu_0(T+\tau)\rangle$ over all initial times $\mathbb{T}_0$—equivalent to fixing $\mathbb{T}_0 = 0$ and integrating over all T—and over all starting abundances $\nu_0(\mathbb{T}_0)$ weighted by the half-Gaussian measure. We work with all times in evolutionary epochs, so that the integral of the distribution of $\nu_0(\mathbb{T}_0)$ for $\nu_0(\mathbb{T}) > 0$ is the rate of successful invasions per unit time, which is 1 in these units.

A self-consistent solution is possible if the result of this calculation gives an exponentially decaying correlation function that can be equated with the correlation function of the cavity drive $C(\tau) = \sqrt{J/b}\ e^{-b\tau}$. Fortunately, this turns out to be the case. An additional self-consistency condition fixes $\sum_i \nu_i = 1$. Put together, these conditions allow us to solve for J and b. We then compute the diversity $\overline{L} = \sum_i \Theta(\nu_i)$. These calculations are carried out in Appendix K. The key results are the self-consistent exponential solution for the correlation function

$$C(\tau) = \frac{2\pi\lambda^2}{Q^2}\exp\left[-\frac{2\lambda}{Q^2}\tau\right] \quad \text{with} \quad \lambda = \frac{1+2\log 2}{4}, \tag{9}$$

and the steady-state diversity given by

$$\overline{L} = \frac{2\log 2}{1+2\log 2}Q^2 \cong 0.58\,Q^2. \tag{10}$$

The timescale of the decay of the correlation function is $O(Q^2) = O(\overline{L})$, consistent with the community turning over every $O(\overline{L})$ successful invasions. We additionally find that $b = 2\lambda/Q^2$ and $J = 4\pi\lambda^3/Q^4$. In Appendix K we verify that our predictions for $C(\tau)$ and $\overline{L}$ agree with numerics.

3. Lifetime and abundance distributions

From the self-consistent solution we can also calculate the distribution of strain *lifetimes*, where the lifetime of a strain i, denoted $\mathcal{T}_i$, is the number of evolutionary epochs for which it has nonzero abundance. In Appendix K we find a distribution

$$p_{\mathcal{T}}(\mathcal{T}) = \frac{2b}{\pi\sqrt{e^{2b\mathcal{T}}-1}}. \tag{11}$$

Figure 5(b) shows good agreement between this analytical result and the numerics: $p_{\mathcal{T}}(\mathcal{T})$ decays exponentially at large $\mathcal{T}$ as expected, and has an integrable divergence as $\mathcal{T} \to 0$, due to the distribution of invaders having a constant density at low abundance (which would be altered if we included the establishment probability coming from demographic fluctuations). These invaders go extinct very quickly, as discussed in Sec. V B. Note that, as successful invaders must balance extinctions, $\langle\mathcal{T}\rangle = \frac{\log 2}{b} = \overline{L}$: the mean diversity is equal to the average strain lifetime in epochs.

The dynamic cavity method also enables us to determine the distribution of strain abundances in steady state, as the solution to a Fokker Planck equation with appropriate boundary conditions and with a source term for the invading strains [Eq. (K12)]. This solution is plotted in Fig. 5(a), in reasonable agreement with the numerical result for the distribution of abundances in the Red Queen phase. As anticipated in Sec. V B, there is a marked difference between the abundance distribution in the eco-evolutionary steady state and the saturated assembled community abundance distribution with $L = \hat{L}$. In the assembled community the abundance distribution is a truncated Gaussian—a half-Gaussian at saturation. By contrast, in the Red Queen steady state, the density of strains vanishes at low abundance due to the constant jostling of the community by invaders which knock low-abundance strains extinct: the effective diffusion coefficient capturing this process is proportional to the fragility.

D. Effects of selective differences

Although we have so far only analyzed the Red Queen phase without selective differences ($\Sigma = 0$), numerics indicate that the phase is robust to $\Sigma > 0$ as well (Appendix E). We now discuss more systematically the effects of a distribution of general fitnesses, focusing first on the distribution of strain lifetimes, $p_{\mathcal{T}}(\mathcal{T})$ and the correlation function $C(\tau)$, both of whose decay at large τ is dominated by strains with large s. We then consider large Σ and give scaling arguments for how the community evolves into the tail of the s distribution, reaching a steady state even in the presence of general fitnesses. These arguments allow us to find the dependence of $\overline{\Upsilon}$ and $\overline{\log p_{\text{inv}}}$ on Σ. When the distribution of extant s_i pushes into the tail of the $p(s)$ distribution, all of the steady-state properties are determined by the form of $p(s \to \infty)$—thus we consider non-Gaussian $p(s)$ with variable rate of decay for large s, and we also comment on the behavior if $p(s)$ has bounded support.

1. Correlation function and lifetime distribution

From the general structure of the cavity analysis, we can attempt to understand the behavior of a strain with anomalously large s (we suppress the strain index for notational convenience). The behavior is simplest in the fully asymmetric LV model, for which, even with $\Sigma > 0$, the cavity drive $\zeta(T)$ still has mean zero and determines the abundance of the cavity strain when in the community [Eq. (6)]. Extinction occurs when $\xi = \zeta + s - \Upsilon \searrow 0$, i.e., $\zeta \searrow \Upsilon - s$ which is unlikely for large s since, for large-s strains, ξ will fluctuate around $s - \overline{\Upsilon}$, staying far from zero for a long time. The rare event for ζ to decrease to $\overline{\Upsilon} - s$ has, quite generally for well-behaved stationary Gaussian processes, probability of order $\sim\exp[-(s-\overline{\Upsilon})^2/2\sigma_\zeta^2]$ for sufficiently large $s - \overline{\Upsilon}$, with the extinction most likely to occur once the correlation has decayed significantly from its initial $C(0) = \sigma_\zeta^2$. However, we need to examine if the cavity *Ansatz* of a well-behaved Gaussian process is consistent when large s can occur.

The effects of large s are most clearly understood by their effects on the average strain lifetime $\langle\mathcal{T}\rangle$, which is equal to $\overline{L}$ in steady state. The naively predicted behavior is that the average lifetime of strains with large s is at least the inverse of their asymptotic extinction rate. Thus, conditioned on s, the mean $\mathcal{T}$ is $\langle\mathcal{T}|s\rangle \sim \exp[(s-\overline{\Upsilon})^2/2\sigma_\zeta^2]$. The unconditional expected lifetime is then

$$\langle\mathcal{T}\rangle = \frac{\int \mathcal{P}(\text{invade}|s)p(s)\langle\mathcal{T}|s\rangle ds}{\int \mathcal{P}(\text{invade}|s)p(s)ds}, \tag{12}$$

where $\mathcal{P}(\text{invade}|s)$ is the invasion probability of a strain with general fitness s, which for large s will be close to unity. The effects of large s on both the correlation function and the lifetime distribution are driven by the fact that long-lived strains have abundance $\nu \sim s - \overline{\Upsilon}$ for most of their lifetime

(with coefficient $1/Q$ for the fully asymmetric LV model). Thus $C(\tau)$ is dominated at long τ by the fraction of strains whose lifetime is at least τ. We therefore expect that, for large τ,

$$C(\tau) \approx \int_{\tau}^{\infty} (\mathcal{T}-\tau)\langle \nu^2|\mathcal{T}\rangle p_{\mathcal{T}}(\mathcal{T})d\mathcal{T}, \tag{13}$$

with the $\mathcal{T}-\tau$ term reflecting the fact that longer-lived strains contribute more to the correlation function, and the ν^2 term likely yielding $\log(\tau)$ factors. If $p(s)$ falls off sufficiently rapidly, the expression for $\langle\mathcal{T}\rangle$ in Eq. (12) is convergent and the correlation function concomitantly decays to zero. To study how the decay of $p(s)$ affects the long-time behavior of $C(\tau)$ and $p_{\mathcal{T}}(\mathcal{T})$, we consider a family of distributions $p(s) \sim e^{-(s/\Sigma)^{\psi}/\psi}$ with the exponent ψ serving as the tuning parameter that controls the rate of decay of the tail of $p(s)$.

For $\psi > 2$, one finds $C(\tau) \sim \exp[-a\log^{\psi/2}(\tau)]$ with $p_{\mathcal{T}}(\mathcal{T})$ decaying similarly so that $\langle\mathcal{T}\rangle < \infty$. The Gaussian case, $\psi = 2$, is marginal with the integral over s convergent only if $p(s)\langle\mathcal{T}|s\rangle$ decays to zero for large s, i.e., for $\Sigma < \Sigma_\times$ where $\Sigma_\times = \sigma_\zeta(\Sigma_\times)$. For $\psi = 2$ and $\Sigma < \Sigma_\times$, $C(\tau)$ decays as a Σ-dependent power of τ.

When $\Sigma > \Sigma_\times$ in the Gaussian case, as well as for any $\psi < 2$, the integrals over s that determine both $\langle\mathcal{T}\rangle$ and the correlation function appear to diverge. Thus this naive analysis, which relied on the Gaussian rarity of large ζ fluctuations, must break down—or the Red Queen steady state must cease to exist. To better understand what happens for $\Sigma > \Sigma_\times$, we turn to numerics.

In Appendix O we explore numerically the effects of Σ for the fully asymmetric LV model with Gaussian $p(s)$ across a range of Σ. In spite of the naive expectation from the above analysis of a loss of steady state for large enough Σ, the numerics show a steady state even for $\Sigma \gg \Sigma_\times \cong 1.4/Q$ (note that for the LV model the natural scale of Σ is $1/Q$ instead of unity for the linearized model). The distributions $p_{\mathcal{T}}(\mathcal{T})$ appear to have a finite first moment and $C(\tau)$ still appears to decay to zero (Fig. 21).

We discuss possible resolutions of this apparent discrepancy between the cavity analysis and the observed behavior in Appendix O, but for now must leave unresolved the subtle effects of longish tails of $p(s)$ and concomitant effects of large deviations from typical. For the present purposes, we note that empirically, even for quite large $\overline{L}$, the primary conclusions of the cavity analysis, here and below, are at least semiquantitatively good.

2. *Evolutionary transients and steady state*

We now analyze the evolutionary dynamics in the limit of large Σ assuming that either $p(s)$ decays sufficiently rapidly for the problems with rare events not to matter, or that the analysis ignoring them is correct for averaged properties. In what follows we mainly discuss the *linearized model* in which the distribution of biases has $O(\sqrt{D/L}) = O(1)$ width—this can be translated to the LV model by inserting appropriate factors of $\sqrt{D}$, which plays the same role as Q.

For $\Sigma \gg 1$, there is a long evolutionary transient during which the distribution of extant s_i creeps into the tail of $p(s)$. The tail of $p(s)$ therefore determines the long-term evolutionary dynamics, including—when it occurs—the properties of the Red Queen phase. As the population mean s (and therefore Υ) increases, the extant s values will be close to Υ with the scale of their variation set by the shape of $p(s)$ near Υ. The width of the extant s_i distribution will be of order

$$\Delta(\Upsilon) \equiv \mathrm{E}\left[s_I - \Upsilon | s_I > \Upsilon\right] \approx -\left[\frac{d\log p(s)}{ds}\right]^{-1}\Bigg|_{s=\Upsilon}, \tag{14}$$

which decreases with increasing Υ for $p(s)$ decaying faster than exponentially and increases with Υ for $p(s)$ decaying slower than exponentially. At any point in the evolution, once $\Upsilon \gg \Sigma$, the distribution of extant s_i is concentrated around Υ and the system behaves as if there were an exponential distribution of the s_i with scale $\Delta(\Upsilon)$.

Given this exponential approximation to the general fitness distribution, it is helpful to understand the behavior of an *assembled* community with an exponential $p(s)$ when the extant s_i are large enough that the lower bound of $p(s)$ does not matter—this will determine a key relationship between L and Υ. If we assemble a large number of strains with an exponential distribution of general fitnesses $p(s) = \frac{1}{\Delta}e^{-s/\Delta}\Theta(s)$, for sufficiently large Δ the ecological cavity analysis (Appendix J) yields an assembled community with $L = \kappa D/\Delta$ in the linearized model. We therefore expect that during the transient evolution, the community will reach a slowly changing quasi-steady-state with diversity given, at any point in time, by $L \sim \kappa D/\Delta(\Upsilon)$.

We have shown above that the behavior of Υ determines the behavior of L. How does Υ depend on T during the evolutionary transient and into the Red Queen steady state? The increase of Υ is tied to the increasing population mean s defined as $\langle s\rangle \equiv \sum_i \nu_i s_i$, so we are led to analyze how $\langle s\rangle$ changes in evolutionary time. Each successful invader I must have $\zeta_I + s_I > \Upsilon$. The probability density of s_I conditioned on invasion is (see Appendix of Ref. [45])

$$p(s_I|\text{invasion}) \approx \frac{1}{\Delta}\exp\left[-\frac{s_I-\Upsilon}{\Delta} - \frac{\sigma_\zeta^2}{2\Delta^2}\right]\Phi\left[\frac{s_I-\Upsilon}{\sigma_\zeta}\right], \tag{15}$$

where $\sigma_\zeta^2 \equiv \mathrm{Var}[\zeta_I] \sim D/L$ and Φ is the CDF of the standard normal distribution. It follows that

$$\mathrm{E}\left[s_I - \Upsilon|\text{invasion}\right] = \Delta(\Upsilon) - \frac{\sigma_\zeta^2}{\Delta(\Upsilon)}. \tag{16}$$

For $p(s) \sim e^{-(s/\Sigma)^{\psi}/\psi}$, we have $\Delta(\Upsilon) \approx \Sigma^{\psi}/\Upsilon^{\psi-1}$. With stretched exponential decay of $p(s)$, $\psi < 1$, $\Delta(\Upsilon)$ is an increasing function—therefore Υ will run away to positive values with L concomitantly decreasing. The diversity thereby crashes and eventually the community will reduce to a single strain that goes extinct when a fitter strain invades.

For $\psi > 1$ (including the Gaussian case $\psi = 2$), initially Δ is of order Σ. From the above analysis, we expect initial diversity $L \sim \kappa D/\Delta$. With Δ large, invasions will tend to increase Υ, thereby decreasing Δ and increasing L. Eventually Δ will decrease to $\overline{\Delta} \sim \sqrt{D/L} \sim 1/\kappa$, independent of ψ, so that the widths of the extant s_i and ζ_i distributions are comparable. Concomitantly, L will saturate at $\overline{L} \sim \kappa D/\overline{\Delta} \sim \kappa^2 D$,

and Υ will saturate at $\overline{\Upsilon} \sim (\Sigma^{\psi}/\overline{\Delta})^{1/(\psi-1)}$, with $\overline{\log p_{\rm inv}} \sim (\Sigma/\overline{\Delta})^{\psi/(\psi-1)}$.

In addition to characterizing the Red Queen steady state, we can analyze the transient evolutionary dynamics that occur when $\Sigma \gg 1$ and $\psi > 1$. For Δ large, as it will be initially, the change in Υ due to a successful invasion by strain I is roughly $\Delta(\Upsilon)/L$, coming from the fitness increment multiplied by the $O(1/L)$ typical abundance ν_I. With $L \sim D/\Delta$, this gives $d\Upsilon/dT \sim \Delta(\Upsilon)^2/D$, where T measures successful invasions (or epochs). On average $d\Upsilon/dT$ is positive until the contribution of ζ_I becomes comparable to $s_I - \Upsilon$, at which point $\Upsilon \approx \overline{\Upsilon}$ and $\Delta \approx \overline{\Delta}$.

The transient behavior is found by integrating up the typical change per invasion: roughly $d\Upsilon/dT \sim \Delta(\Upsilon)^2/D \sim \Sigma^{2\psi}/D\Upsilon^{2\psi-2}$ until Υ is close to $\overline{\Upsilon}$. Thus during the transient

$$\Upsilon \sim \left(\frac{\Sigma^{2\psi} T}{D}\right)^{\frac{1}{2\psi-1}}, \tag{17}$$

$$\frac{L}{D} \sim \Sigma^{-\frac{\psi}{2\psi-1}} \left(\frac{T}{D}\right)^{\frac{\psi-1}{2\psi-1}}, \tag{18}$$

for a duration

$$T_{\rm transient} \sim D\Sigma^{\psi/(\psi-1)}, \tag{19}$$

after which the steady state ensues.

An important observation is that some aspects of the Red Queen steady state for large Σ are *independent of details* of the s distribution and of Σ: in particular, for large Σ, $\overline{L}$ is independent of Σ, as seen in Appendix E for Gaussian $p(s)$. A key feature of the Red Queen phase is that $\overline{\Upsilon}$ becomes large enough that the distribution of the extant s_i has $O(1)$ width, irrespective of Σ. With the widths of the extant s_i and ζ_i distributions comparable, Υ no longer pushes further into the tail of $p(s)$ since having large s_i is no longer sufficient to guarantee having large invasion fitness.

Figure 6(a) shows the evolution of the extant distribution of general fitnesses for large Σ with Gaussian $p(s)$: the mean increases and the variance decreases as the community pushes out into the tail of $p(s)$. In Fig. 6(b) we plot the dynamics of the width of the extant s_i distribution for various Σ. With time rescaled by the transient timescale $\sim\Sigma^2$ for $\psi = 2$ [Eq. (19)] the curves collapse and we confirm that in the Red Queen steady state, the width of the extant s_i distribution converges to an $O(1)$ value independent of Σ.

In Appendix N we check the predicted transient behavior against numerics for a Gaussian s distribution, for which $\Delta(\Upsilon) \approx \Sigma^2/\Upsilon$, and we have $\overline{\Upsilon} \sim \Sigma^2$ and $-\overline{\log p_{\rm inv}} \sim \Sigma^2$.

During the long initial transient, our scaling arguments say that $\Upsilon \sim T^{1/3}$ and $L \sim T^{1/3}$, with the transient lasting for a number of epochs $T_{\rm transient} \sim D\Sigma^2$, consistent with the rapid growth of transient length with Σ in Appendix E. Perhaps surprisingly, given the potential problems caused by large s_i discussed earlier in this section, the predicted transients and steady-state behavior in both Appendix N and in Fig. 6 seem to persist for Gaussian $p(s)$ even when $\Sigma \gg \Sigma_\times$.

3. Invasion probabilities

Because of large κ-dependent coefficients, even moderate values of Σ can make $p_{\rm inv}$ unrealistically small at steady state.

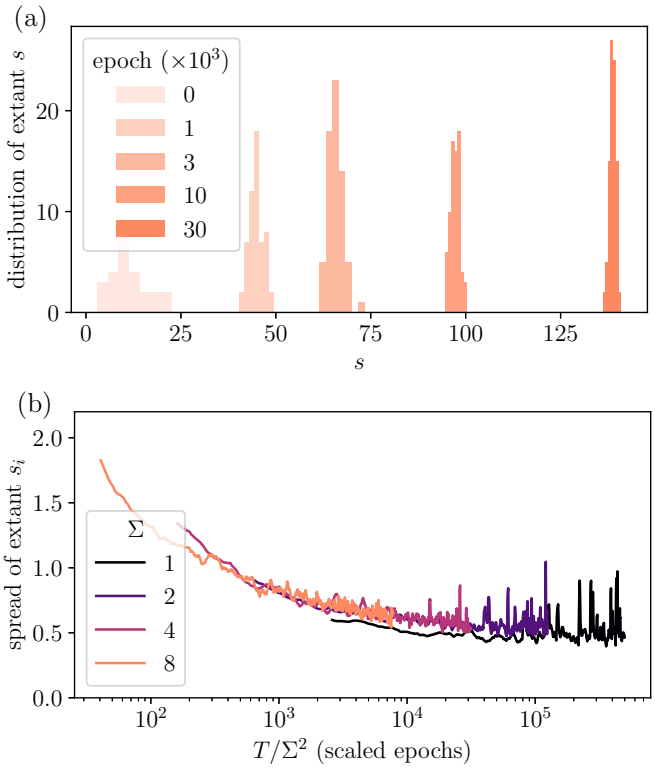


FIG. 6. Distribution of extant general fitnesses, $\{s_i\}$, showing evolutionary transient pushing into the tail of $p(s)$ and decreasing the width of the distribution, with steady-state width independent of Σ. Linear model with $D = 200$, $\kappa = 0.8$, and $\psi = 2$. (a) Distributions for large Σ (here $\Sigma = 8$) showing evolution to larger mean and smaller variance. The last histogram (3×10^4 epochs) is still far from steady state. (Distributions normalized to L). (b) Moving average of standard deviation of s_i with collapse of curves showing approach to an $O(1)$ value independent of Σ. Transient timescale scales with Σ^2 [Eq. (19)]. (Again, for $\Sigma = 8$, the steady state is not reached.) Large fluctuations in the width of the distribution in steady state are caused by anomalously small-s strains invading; these are accompanied by temporary decreases in L (not shown).

For $\Sigma = 1$, $\rho = 0$ (unrelated invaders) and $\kappa = 0.8$, the invasion probability at steady state is less than 10^{-15} (Appendix E) which suggests that $\overline{\Delta}$, which determines $\overline{\Upsilon}$ via $\overline{\Upsilon} \sim \Sigma^2/\overline{\Delta}$, has a value $\overline{\Delta} \cong 0.12$ for $\kappa = 0.8$, substantially smaller than the value for a saturated top-down assembled community, $\hat{\Delta} = (1-\kappa^2)/\kappa = 0.45$ [Eq. (5a)].

In the CR model, invasion probabilities are even smaller: as discussed earlier, the width of the general fitness terms dominates over that of the interactions by a factor of $\sqrt{D}$, corresponding to $\Sigma \sim \sqrt{D}$ in the linearized model. Therefore, for our Gaussian ensemble, the typical $p_{\rm inv}$ in the steady state is exponentially small in D. Indeed we see [Fig. 2(b)] that for $\kappa = 0.4$ and $D = 50$ in the CR model, the invasion probability in steady state is $\approx 10^{-400}$, where for the nonlinear CR model, the distribution of general fitnesses U_i inherits its Gaussian distribution from the Gaussian ensemble of the $\mathcal{G}$ matrix that we studied.

Generally, distributions with faster-decaying tails but similar Σ have higher invasion probability at steady state. However, even in the limit $\psi \to \infty$ which corresponds to a

double exponential decay $p(s) \sim \exp[-\exp(s/\Sigma)]$, the typical steady state $p_{\rm inv}$ still decays exponentially in Σ and therefore as a stretched exponential in D for the CR model.

For distributions that have a sharp cutoff at some maximum s, the invasion probability in steady state is not so small, even with a broad distribution of s. If $p(s) \sim (\Sigma - s)^{\lambda}\Theta(s)\Theta(\Sigma - s)$ for large Σ and some power λ, then the width of the extant s_i distribution scales as $\frac{1}{\lambda+2}(\Sigma - \Upsilon)$. The Red Queen phase obtains when the distribution of extant s has $O(1)$ width, in which case $p_{\rm inv} \sim 1/\Sigma^{\lambda+1}$ in steady state, which decays with Σ much more slowly than in the case of $p(s) \sim e^{-(s/\Sigma)^{\psi}/\psi}$. Importantly, once the Red Queen steady state is established, the population mean $\langle s\rangle$ does not creep ever closer to the cutoff at the maximum s. This is in contrast to the situation with $\kappa = 1$, where Υ would inch up as $\langle s\rangle$ increases towards the maximum s, since both Υ and the Lyapunov function $\mathcal{L} = \Upsilon + \langle s\rangle$ contain a part proportional to $\langle s\rangle$.

We have found that there is a wide range of scalings of $\overline{\Upsilon}$ and the typical steady state $p_{\rm inv}$ as a function of Σ in the linearized (or LV) model, and as a function of D in the CR model. These scalings can be reasonably obtained with different distributions of the s_i or the elements of $\mathcal{G}$ and $\mathcal{F}$. In our numerics we have used Gaussian distributions since they allow us to easily draw invaders conditioned on having positive invasion fitness and to explore issues of *principle* regarding the Red Queen phase. However, quantitative aspects such as the scaling of the invasion probability with D are very nonuniversal. In biological systems, situations where the evolution pushes far into the tail of $p(s)$ are anyway likely dominated by effects that we do not account for, such as the discreteness of genomes.

VI. DISCUSSION

In this work we have explored evolution in a range of ecological models that exhibit diverse stable communities and found that, generically, after a transient period of evolution, the community converges to a Red Queen state of constant turnover with a roughly constant probability of new strains successfully invading. In contrast with previous work which finds continual evolution in selection-driven host-pathogen models [27,55,56], the diverse communities we find here stably coexist and do not rely on continual mutations to replenish the diversity. Here, there are no "kill-the-winner" dynamics in which one subpopulation "chases" the other through trait space; nonetheless constant turnover ensues. The evolving ecology acts as a concrete manifestation of the environmental feedback that can engender continual Red Queen evolution in random landscape models without diversity [18]. The consumer resource model and a linearized approximation to it display a Red Queen phase as long as the consumption and growth are not perfectly correlated ($\kappa < 1$), while the oft-studied perfectly correlated case ($\kappa = 1$) behaves differently and has no Red Queen phase.

Importantly, the Red Queen phase persists in the presence of arbitrarily large general fitness differences s_i provided that their distribution falls off at least exponentially for large s (with a more rapid decay possibly required in the limit of large $\overline{L}$ due to the issues discussed in Sec. V D). Previous analyses of high-dimensional ecological models [20,38] have, as noted earlier, first assumed that in the large-D (or similar) limit, the distribution of general fitnesses (or similar quantities) has a width that is smaller than that of the interactions by a factor of $1/\sqrt{D}$, and then taken a well-behaved large-D limit. We have shown that this assumption is not needed: In a long-evolved community, the width of the distribution of general fitness differences *within the evolved community* will become comparable to the variations between strains of the effects of ecological interactions. Even starting with a parametrically (in D) broad distribution of general fitnesses, some properties (such as the diversity) of evolved communities are roughly independent of the distribution, although other properties (such as the invasion probability of independent strains) can be strongly affected.

If the distribution of the s_i falls off with a large characteristic scale—which occurs naturally in the CR model—there is a long transient period during which evolution is dominated by increases in $\langle s\rangle$, and the diversity gradually increases as the population moves out into the tail of $p(s)$. The transient behavior converges to a steady state when the strain-to-strain variations of s_i and the ecological influences of other strains in the community are comparable—here the diversity of the Red Queen steady state becomes independent of $p(s)$, but the invasion probability at steady state can be extremely small. However, if there is a maximum s, the invasion probability, and hence the community turnover rate, stabilizes at a value that need not be as small.

Below, we discuss some unresolved issues from this current work, and suggest future directions.

A. Other eco-evolutionary phases in well-mixed systems

What other long-term states occur naturally and robustly in eco-evolutionary models with no spatial structure or externally varying environment? For models with perfect symmetry, such as the $\kappa = 1$ CR and linearized models, the presence of a Lyapunov function means that the community can continue to evolve and turn over but will do so more and more slowly. Does such slowing evolution occur away from special slices of parameter space? In the CR and linearized models with a convex Lyapunov function, the $\kappa = 1$ behavior appears to be approached continuously as $\kappa \to 1$, with the Red Queen phase giving way to a phase of monotonic evolution only at $\kappa = 1$. Is the approach towards perfect symmetry similarly smooth in other models?

To answer this question we study evolution in the LV model (via independent invaders) with a tunable symmetry parameter $\gamma \in [-1, 1]$, which quantifies the cross-diagonal correlations of the interaction matrix: $\mathrm{E}[V_{ij}V_{ji}] = \gamma\,\mathrm{E}[V_{ij}^2]$ for $i \neq j$. The diagonal elements of V have mean $-Q$ and variance $1 + \gamma$, while off-diagonals have zero mean and unit variance (Appendix F). This model is a generalization of the totally asymmetric LV model analyzed in Sec. V C. As is the case in the linearized model, $\Upsilon + \langle s\rangle$ is a Lyapunov function of this model when $\gamma = 1$. However, the eco-evolutionary behavior of this model as $\gamma \to 1$ is markedly different from that in the CR or linearized models as $\kappa \to 1$.

In the LV model we see a sharp (in the limit of large Q) qualitative change in the dynamics as a function of γ, around $\gamma_c \approx 0.65$ for $\Sigma = 0$. For $\gamma < \gamma_c$, the Red Queen phase

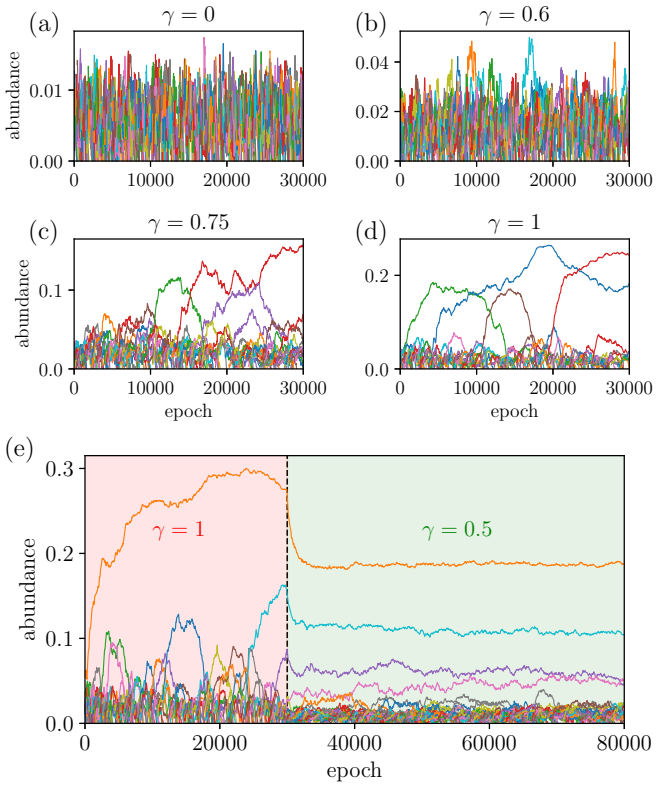


FIG. 7. Strain trajectories (subset of strains shown) for LV model with (a), (b) various γ in Red Queen and (c), (d) oligarch phases. In the oligarch phase, an $O(1)$ fraction of the total abundance aggregates into a few strains which continue to turn over more and more slowly. In panel (e), initially $\gamma = 1$, but after 3×10^4 epochs (vertical dashed line) γ is "quenched" to $\gamma = 0.5 < \gamma_c$. System does not revert back to the original steady state, instead showing several long-lived strains persisting. $Q = 20$ and $\Sigma = 0$ throughout.

obtains and there is continual turnover with the mean general fitness Υ and the invasion probability p_{inv} roughly constant. However, for $\gamma > \gamma_c$, we find no steady state. This transition is not apparent in the dynamics of the diversity L, which is still roughly steady over time, but shows up strikingly in other properties.

For $\gamma > \gamma_c$, the dynamics continually slows down, with Υ ever increasing and p_{inv} thereby decreasing so that the community becomes progressively harder to invade. This behavior, which we have dubbed an *oligarch* phase, has surprising features that we do not yet understand. Most strikingly, the evolutionary dynamics show that a few strains—the oligarchs—become far more abundant than the others.

Figures 7(a)–7(d) shows the abundance dynamics in both the Red Queen and oligarch phases, with the corresponding trajectories of L, Υ, and p_{inv} shown in Appendix H. For $\gamma = 0$ and 0.6, the constant turnover of the Red Queen phase is evident in the narrow distribution of strain abundances and lifetimes. However, for $\gamma = 0.75$ and 1, an $O(1)$ fraction of abundance "condenses" into an $O(1)$ number of oligarch strains which still slowly turn over.

Qualitatively, the oligarch phase appears to occur due to feedback between the abundant strains and the invaders for larger γ: if a prospective invader I is to successfully enter the community it must have a positive contribution to its invasion fitness from an abundant strain ℓ, i.e., $V_{I\ell} > 0$. Then, since γ is strongly positive, the invader will tend to have $V_{\ell I} > 0$ as well, thus increasing the abundance of the oligarch strain ℓ. However, this argument does not explain the turnover of oligarchs, nor their absence in the linearized model. The non-steady-state nature of the oligarch phase makes analysis difficult since correlation and response functions of strain abundances in evolutionary time depend on two times, not just time differences.

The transition between the Red Queen and oligarch phases (discussed in Appendix H) appears to be first-order-like. As then expected, it persists for (at least) small general fitness variations, with γ_c depending on Σ in a way that we have not investigated. Hysteresis near the transition is observed when γ is suddenly reduced through γ_c (Fig. 7(e) and Appendix H). Strikingly, if one evolves with $\gamma > \gamma_c$ for some time, and then rapidly quenches to $\gamma < \gamma_c$, one can obtain a steady state which exhibits aspects of both the Red Queen and oligarch phases: a handful of high-abundance strains persists at roughly constant abundance without turning over, while the majority of strains turn over steadily as they would in the "normal" Red Queen phase. In this coexisting state between the two phases, Υ and p_{inv} are roughly constant although remain respectively elevated and suppressed relative to their values in the Red Queen phase, while L does not show such a signature (Appendix H). Further analysis of such a steady—and therefore more analytically tractable—oligarch state could lead to a better understanding of the oligarch phase as well as the transition between the two eco-evolutionary phases.

What classes of models can exhibit an oligarch phase? Naively, the heuristic argument above would suggest an oligarch phase in the linearized model, although such a phase does not seem to exist. This difference in behaviors is likely related to the properties of the totally symmetric LV and linearized models: the symmetric linearized model always has a *convex* Lyapunov function whereas the LV Lyapunov function becomes nonconvex for sufficiently large S [39]. What other phenotype models—perhaps with more general interaction via chemicals—can display an oligarch phase of eco-evolutionary dynamics? Given aforementioned subtleties regarding the tails of the general fitness distribution and the long-lived strains with large s_i, one must also ask how the tails of the distribution of the elements of V affect the oligarch phase.

If some fraction of the D resource-mediating chemicals have negative correlations between the corresponding columns of G and F, there exists a $D \to \infty$ limit of the model with the fraction of negatively correlated resources carefully tuned so that the diagonal terms in the interaction matrix $\mathcal{V} = -GF^\intercal$ are larger than the off-diagonal by a factor analogous to the LV model's Q, with the cross-diagonal correlations across the diagonal of the interaction matrix remaining of order unity. In such a case, as $D \to \infty$ one recovers a LV model with $\gamma \neq 0$. We have seen numerically that such a linearized consumer resource model with some fraction of negatively correlated resources displays an oligarch phase even for finite D. It therefore seems likely that an appropriately tuned nonlinear consumer resource model would also display an oligarch phase, although perhaps only after a long transient due to the broad distribution of general fitnesses. In any case, the

presence of the oligarch phase in the LV model illustrates the range of phenomenology that can occur in different models, with each distinct eco-evolutionary phase robust over a range of parameters.

B. Parent-mutant correlations

In this work, we have primarily studied a process that is not, strictly speaking, evolution: new strains that are introduced are unrelated to extant strains. Nevertheless, this process is crucially different from well-studied models of top-down assembly. In our pseudo-evolutionary process, strain extinctions are permanent, which qualitatively changes the shape of the abundance distribution: in steady state few strains with low abundance exist as these are most easily driven extinct by invasion-induced perturbations.

We have shown that introducing mutant strains related to extant parent strains (with correlation $\rho > 0$) results in very similar behavior if the mutant and parent are not strongly correlated phenotypically.—ρ not too close to unity. However, if the effects of mutations are small—surely relevant biologically [57]—diversification is slower because of the increased importance of general fitness differences. This regime raises important questions: How do the existence and properties of a diverse Red Queen phase depend on the magnitude of the mutational effects when these are very small—i.e., in the limit $\rho \nearrow 1$? What happens with a distribution of mutational effect sizes? What is the interplay between general fitnesses and parent-mutant correlations? What types of phylogenies are generated by the Red Queen eco-evolutionary process? Do they resemble any known class of coalescent trees? We address these issues in a future paper.

C. Changing environments

The CR and linearized resource models naturally allow study of how evolving communities are affected by *extrinsic changes* in the environment. During the initial transient period of evolution before the Red Queen phase obtains, the community in the CR model adapts to its resource supply by aligning the growth vectors of strains (rows of $\mathcal{G}$) with the resource supply vector $\mathbf{K}$. If the supply vector $\mathbf{K}$ changes extrinsically after this initial period of evolution, the distribution of U_i will broaden and the diversity will crash, only to rebuild as the community readapts to the new resource supply vector. This scenario can be studied in the linearized model if we parametrize the general fitnesses through an interaction with the environment, and define the resource vector as $R_\alpha = k_\alpha - \sum_j F_{j\alpha}\nu_j$, with the vector $\mathbf{k}$ proportional to deviations of the resource supplies from some reference values. This parametrization gives rise to selective differences $s_i = \sum_\alpha G_{i\alpha}k_\alpha$: therefore changing $\mathbf{k}$ will change the s_i.

What do the eco-evolutionary dynamics look like for a continually changing supply vector, either driven extrinsically or through some further feedback from the ecological community? How can this externally driven continual evolution be distinguished from the self-driven Red Queen evolutionary phase?

D. Clonal interference and sexual reproduction

In this work we have taken advantage of a separation of timescales between evolution and ecology. However, such a separation is not at all guaranteed to exist, particularly in host-pathogen systems [58], in large microbial populations with abundant beneficial mutations [9], or when mutational effects are small enough that ecological timescales are slow. Investigation of the dynamics of ecological models in a "clonal interference" regime, with multiple ecological mutants emerging before any single one of them has reached its fixed-point abundance, is an open challenge of definite relevance to microbial populations.

In addition, phenotype models such as the consumer resource or linearized models accommodate a natural parametrization of sexual or quasisexual reproduction between different strains, which could create new phenotypes that are a mix of two (or more) extant phenotypes. Such models could therefore be fertile ground for understanding the interplay of diverse ecology with horizontal gene transfer, which is surely relevant for natural microbial populations.

E. Spatial structure

Spatial structure and transport remains an understudied but crucial feature of ecological and evolutionary models. Gut and skin microbiota associated with different individuals are realizations of distinct ecological communities coupled by person-to-person migration, with new mutants arising locally and having the potential to spread across many different people [59,60]. We have found that even without spatial structure, the eco-evolutionary behavior of stable communities is quite rich. However, generalizing our analyses to spatially mean-field "island" models with all-to-all migration—roughly like microbes from person-to-person—is a very interesting direction, as are extensions to finite spatial dimensions.

If the potential number of niches (roughly D or Q^2 in our models) is not large enough for all the strains in the system to coexist locally, a larger number of distinct strains could still inhabit different spatial locations. Is this heterogeneity maintained when strains migrate between spatial patches or does its persistence require evolution to continually produces new variants? How much diversity can be stably maintained, and how does this grow with the spatial size or number of microbiome hosts?

F. Other spatiotemporal eco-evolutionary phases

The range of models we have studied in this work hints at the generality of the Red Queen phase in ecological models with stable communities. However, with spatial structure there is another ecologically diverse phase known: a state of spatiotemporal ecological chaos in which the spatial structure stabilizes the fluctuating ecological dynamics against global extinctions, leading to an ecological state that is exponentially long-lived in its spatial size [24]. In previous work [45] we studied gradual evolution and assembly in such a spatiotemporally chaotic population. In the simplest predator-prey-like models that yield this phase, there are no niches created by different resources: the model corresponds to a $Q = 0$ Lotka-Volterra model with negative cross-diagonal correlations ($\gamma < 0$). However, the dynamical nature of the ecological phase allows one to pack unlimited diversity into a single niche, even when strains are added gradually through evolution. In the presence of general fitness differences,

diversification to arbitrarily large L can still occur, but the evolution continually slows down, with p_{inv} becoming smaller and smaller as the general fitnesses (and hence Υ) continue to increase.

We can understand the tendency of these host-pathogen-like models towards unlimited diversity from the phase diagram of the top-down assembled community. If one starts in the unstable phase, then increasing the diversity does not bring one closer to the phase boundary between stable and unstable phases; in fact it takes one further from this boundary. There is no natural diversity scale and no divergent fragility that would limit the community's diversity. By contrast, for the stable dynamics studied here, the presence of sufficiently strong niche interactions sets a scale for the diversity which is $\overline{L} \sim D$ in the CR and linearized models, and $\overline{L} \sim Q^2$ in the LV model. This niche structure qualitatively changes the outcome of the evolutionary dynamics. Nevertheless, one can ask: could a spatiotemporally chaotic phase exist—and perhaps enable more persistent diversity—in resource models?

G. Diversity across phylogenetic scales

Experimental data have shown that strain dynamics can differ based on the resolution at which they are observed, with closely related strains undergoing larger fluctuations than family level groupings [57,61]. Further theoretical studies of ecological dynamics across scales of diversity is needed. How—hopefully already in simple models—can evolution create diversity over many phylogenetic scales, and how does this diversity interact with the ecological dynamics across such scales? Can evolution robustly lead to ecologically chaotic dynamics of microdiversity but stable dynamics at coarser levels of differentiation?

H. Prospects

The models and scenarios described in this work provide a simple framework and examples in which ecology produces an evolutionary Red Queen phase. The analytical tractability of the models in the simplest case—and the generality and robustness of the Red Queen phase we have found—gives hope that the study of simple models might advance our understanding of the evolution and maintenance of fine-scale biodiversity in natural microbial communities.

ACKNOWLEDGMENTS

We thank Pankaj Mehta and Gabriel Birzu for illuminating discussions. This research was supported in part by the National Science Foundation (NSF) via NSF PHY-2210386 to DSF. A.M. and D.S.F. acknowledge support from the Kavli Institute for Theoretical Physics (KITP) through the Gordon and Betty Moore Foundation Grants No. 2919.02, No. NSF PHY-1748958 and No. PHY-2309135, the Heising-Simons Foundation, and the Simons Foundation (216179, LB). We thank the Stanford Sherlock cluster for computing resources.

DATA AVAILABILITY

Code and output data available upon request.

APPENDIX A: STRUCTURE OF CONSUMER RESOURCE MODEL

1. Lyapunov function

The Lyapunov function of the symmetric CR model with $\mathcal{G} = \mathcal{F}$ is

$$\Lambda = \sum_\alpha K_\alpha \log \frac{K_\alpha}{\mathcal{R}_\alpha} - m \sum_i n_i, \tag{A1a}$$

$$\text{with} \quad \mathcal{R}_\alpha = \frac{K_\alpha}{\omega + \sum_{j=1}^S \mathcal{F}_{j\alpha} n_j}. \tag{A1b}$$

Λ is a convex function of the $\{n_i\}$ and increases under the ecological dynamics of Eq. (1) (if resource dynamics are fast). Therefore, Λ also increases under the evolutionary dynamics even with extinctions, as extinct strains do not contribute [37]. Since the same fixed point is found regardless of the resource dynamical timescale, it does not matter that Λ is a Lyapunov function for the effective consumer dynamics when the resources are integrated out: it is still monotonic from fixed point to fixed point as strains are introduced.

2. Nondimensionalizing

With the choice $\omega = 0$, we define $\tilde{\mathcal{R}}_\alpha = \mathcal{R}_\alpha \mu_g / d$, $\tilde{\sigma}_g = \sigma_g/\mu_g$, $\tilde{\mathcal{G}} = \mathcal{G}/\mu_g$, $\tilde{\sigma}_f = \sigma_f/\mu_f$, $\tilde{\mathcal{F}} = \mathcal{F}/\mu_f$, and $\tilde{K}_\alpha = \mu_g K_\alpha / d\mu_f$, such that Eq. (1) reads

$$\frac{dn_i}{d(mt)} = n_i \left(\sum_\beta \tilde{\mathcal{G}}_{i\beta} \tilde{\mathcal{R}}_\beta - 1 \right), \tag{A2a}$$

$$\frac{d\tilde{\mathcal{R}}_\alpha}{d(\mu_f t)} = \tilde{K}_\alpha - \tilde{\mathcal{R}}_\alpha \sum_j \tilde{\mathcal{F}}_{j\alpha} n_j. \tag{A2b}$$

Therefore, κ, D, $\tilde{\sigma}_g$, $\tilde{\sigma}_f$ and the rescaled supply vector $\tilde{\mathbf{K}}$ are the only quantities that matter for the fixed-point abundances. We set $\mu_f = \mu_g = 1$ and $d = 1$ without loss of generality. The fixed-point abundances $\{n_j^*\}$ and $\{\mathcal{R}_\alpha^*\}$ satisfy the equations

$$\sum_\beta \mathcal{G}_{i\beta} \mathcal{R}_\beta^* = 1 \; \forall \; i \in \mathcal{C}, \quad \mathcal{R}_\alpha^* = \frac{K_\alpha}{\sum_j \mathcal{F}_{j\alpha} n_j^*}. \tag{A3}$$

We do not expect that the variance of the components of $\mathbf{K}$ plays an important role, so we set $K_\alpha = K$. The value of K matters only insofar as it determines the scale of the $\{n_j^*\}$, which can be rescaled to eliminate K.

APPENDIX B: PARENT-MUTANT CORRELATIONS IN CONSUMER RESOURCE MODELS

We can draw mutant phenotypes to be correlated with the phenotype of its parent, which is chosen with probability proportional to its abundance in that epoch. Given the parent phenotype—the vectors $\boldsymbol{\mathcal{G}}_P$ and $\boldsymbol{\mathcal{F}}_P$—we draw a mutant phenotype $\boldsymbol{\mathcal{G}}_M$ and $\boldsymbol{\mathcal{F}}_M$ which has elementwise Pearson correlation ρ with the parent phenotype. We also enforce elementwise correlation κ between $\boldsymbol{\mathcal{G}}_M$ and $\boldsymbol{\mathcal{F}}_M$. In Fig. 8 we show numerical results from this process. The trends are similar to what we observe when we vary ρ in the linearized model and, importantly, show that there is still a Red Queen steady state across a range of ρ, with the phase qualitatively

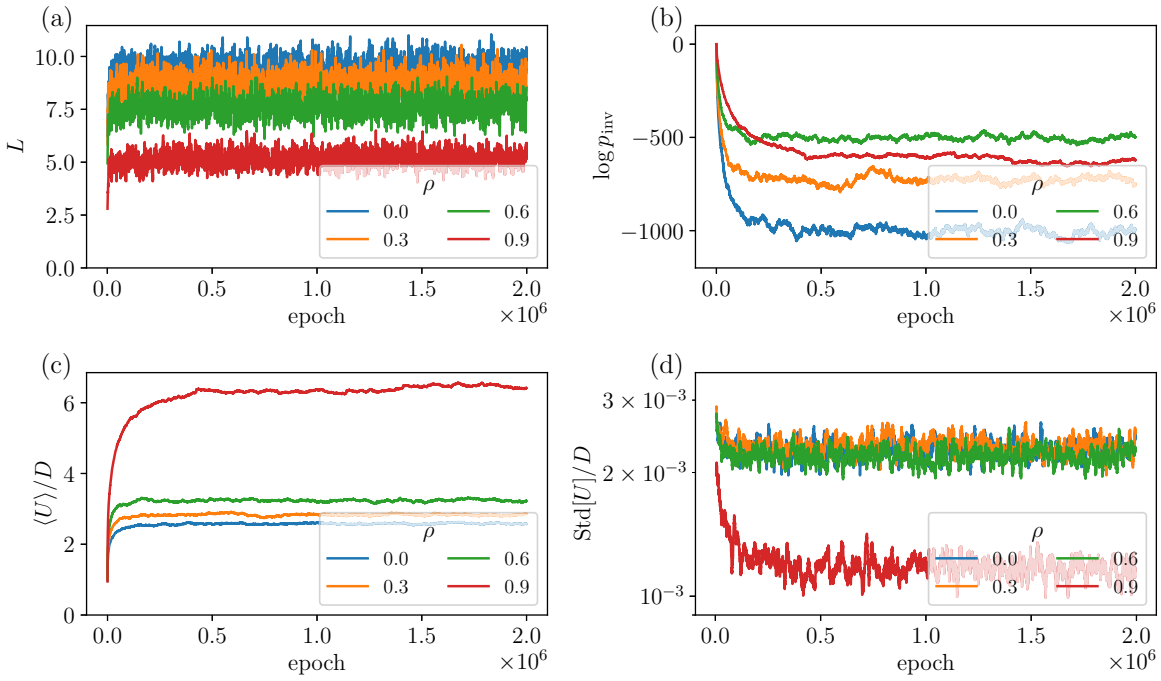


FIG. 8. Effects of parent-mutant correlation, ρ, on evolution in CR model with $D = 50$, $\kappa = 0.4$. (a) Steady state $\overline{L}$ is a decreasing function of ρ. (b) $\overline{\log p_{\rm inv}}$ is a nonmonotonic function of ρ, largest for intermediate ρ. (c) After initial transient, community-mean general fitness $\langle U\rangle$ saturates at a steady-state value. $\overline{\langle U\rangle}$ increases with ρ as distribution of extant U_i pushes farther into tail. (d) $\overline{{\rm Std}[U]}$, the steady-state standard deviation of the extant U_i distribution, decreases with ρ.

similar to the $\rho = 0$ phase studied in depth. Of particular phenomenological interest is the dependence of $\overline{\log p_{\rm inv}}$ on ρ, which displays a nonmonotonicity. We deliberately stay away from the limit $\rho \to 1$, which is left for future work as additional subtleties arise.

APPENDIX C: RELATION BETWEEN CONSUMER RESOURCE AND LINEARIZED MODELS

We here connect the CR model to the linearized model, which we have argued captures the behavior of the nonlinear CR model and is easier to analyze. The linearization of $\boldsymbol{\mathcal{R}}^*$ as a function of the $\{n_j^*\}$ is justified when the mean of $\sum_j \mathcal{F}_{j\alpha} n_j^*$ over α (weighted by the K_α) is larger than its variations. In Fig. 9(a) we show that this condition is met in the Red Queen steady state for $\kappa = 0.4$. Figure 9(b) shows that the total consumer population is roughly constant in evolutionary time, which motivates a description in terms of fractional abundances for the linearized model. We carry out an expansion which allows us to separate out general fitness and interaction contributions to each strain's growth rate.

In the Red Queen steady state the means of $\mathcal{F}$ and $\mathcal{G}$ are greater than 1 due to evolutionary conditioning—for example in Fig. 3(a), the plot for $\kappa = 0.4$ shows that the mean of $\mathcal{G}$ plateaus at $\langle U\rangle/D \approx 2.5$. The mean of $\mathcal{F}$ is smaller than that of $\mathcal{G}$ but still larger than 1. If we define $\langle \mathcal{F}\rangle \equiv \frac{1}{NK_{\rm tot}} \sum_{j,\beta} \mathcal{F}_{j\beta} n_j^* K_\beta$ and $f_{i\alpha} = \mathcal{F}_{i\alpha} - \langle \mathcal{F}\rangle$, then we have $\sum_j \mathcal{F}_{j\alpha} n_j^* = \langle \mathcal{F}\rangle N + \sum_j f_{j\alpha} n_j^*$. Defining $\mathcal{N} = \langle \mathcal{F}\rangle N$, we can Taylor expand the resource abundances, yielding

$$\mathcal{R}_\alpha^* \approx \frac{K_\alpha}{\mathcal{N}} \left(1 - \frac{\sum_j f_{j\alpha} n_j^*}{\mathcal{N}}\right), \tag{C1}$$

where our definition of $\langle \mathcal{F}\rangle$ ensures that $\frac{1}{D}\sum_\alpha \mathcal{R}_\alpha^* = K_{\rm tot}/\mathcal{N}D$. If we approximate this expansion as an equality, at the fixed point the condition on the extant strains that their growth rates vanish becomes

$$0 = \sum_\beta \mathcal{G}_{i\beta} \frac{K_\beta}{\mathcal{N}} \left(1 - \frac{\sum_j f_{j\beta} n_j^*}{\mathcal{N}}\right) - 1 \quad \text{for } i \in \mathcal{C} \tag{C2a}$$

$$\Rightarrow 0 = \frac{1}{\mathcal{N}} \sum_\beta g_{i\beta} K_\beta - \frac{1}{\mathcal{N}^2} \sum_{\beta,j} g_{i\beta} K_\beta f_{j\beta} n_j^* - \frac{\langle \mathcal{G}\rangle}{\mathcal{N}^2} \sum_{\beta,j} K_\beta f_{j\beta} n_j^* + \frac{\langle \mathcal{G}\rangle}{\mathcal{N}} \sum_\beta K_\beta - 1 \quad \text{for } i \in \mathcal{C}, \tag{C2b}$$

where we have similarly defined $\langle \mathcal{G}\rangle \equiv \frac{1}{NK_{\rm tot}} \sum_{i,\beta} \mathcal{G}_{i\beta} n_i^* K_\beta$ and $g_{i\beta} = \mathcal{G}_{i\beta} - \langle \mathcal{G}\rangle$ so that the average of $\sum_\beta g_{i\beta} K_\beta$ over i (weighted by n_i^*) vanishes. This form allows us to separate out the general fitness and interaction pieces in the growth rate of each strain. On the right-hand side of Eq. (C2b), the first term contains the general fitnesses and is proportional to $U_i - \langle U\rangle$. The second term is the contribution to the growth rate of strain i from the ecological interactions, and the last three terms are i-independent and constitute an effective Lagrange multiplier that keeps N roughly constant when $\langle \mathcal{G}\rangle \gtrsim O(1)$.

In the assembled community, $\mathcal{N} = O(L)$ and the $f_{j\alpha}$ and $g_{i\beta}$ have width σ_f and σ_g, respectively. Therefore, the width of the general fitness term scales as $\sigma_g\sqrt{D}/L$ and the width of the interaction piece scales as $\sigma_f\sigma_g\sqrt{LD}/L^2$. Since $L = O(D)$ in an assembled community for sufficiently large S, with each of the n_j^* being $O(1)$, the distribution of general fitnesses is a

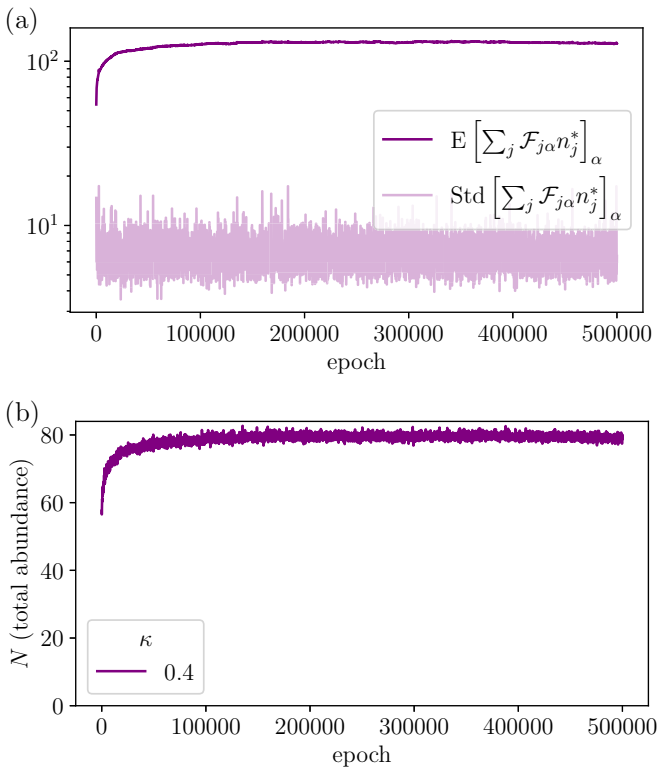


FIG. 9. In the Red Queen phase of the CR model ($D = 50$, $\kappa = 0.4$), linearization of resource abundances is justified, and total consumer population is roughly constant. (a) At steady state, consumption rates of resources, $\sum_j \mathcal{F}_{j\alpha} n_j^*$, have a narrow distribution, justifying expanding variations in this quantity out of the denominator. $K_\alpha = 1\ \forall\ \alpha$, so weighting means and standard deviations by K_α does not affect results. (b) The total consumer population at the fixed point, $N = \sum_j n_j^*$, is roughly constant in the Red Queen phase, motivating switching to fractional abundances.

factor of $\sqrt{D}/\sigma_f$ wider than the distribution of drives. Since σ_f and σ_g are $\lesssim 1$ in order to make the majority of entries in $\mathcal{F}$ and $\mathcal{G}$ positive, the general fitnesses dominate in the assembled community. However, when the community evolves, the $\boldsymbol{\mathcal{G}}_i$ vectors elongate in the direction of the supply vector $\mathbf{K}$. As they push into their Gaussian tail, they decrease the width of the distribution of extant $\sum_\beta g_{i\beta} K_\beta / \mathcal{N}$, which is governed by the variations of the $\mathbf{g}_i$ vectors in the direction parallel to $\mathbf{K}$. But the width of the extant drive distribution does not decrease and this allows the distributions of general fitnesses and growth rates from ecological interactions to attain comparable widths in an evolved community with $L = O(D)$ (Sec. III B).

APPENDIX D: ALGORITHM TO FIND FIXED POINT

Integrating the ecological dynamics to find their stable fixed point is too slow to observe many of the features of the long-term evolution of interest. To probe longer evolutionary trajectories, we have used a heuristic method to find the fixed point; we then check that it is stable and cannot be invaded by any of the extinct strains. Below we describe the algorithm for the linearized and LV models, and a modification of the algorithm for the CR model. For the LV and linearized models, given an interaction matrix, the procedure can be modified to include selective differences by adding s_i to the ith row of V: for ease of notation we refer to this effective interaction matrix simply as V. The main challenge is to find the set $\mathcal{C}$ of strains in the community that has nonzero abundance at the stable uninvadable fixed point. We accomplish this with an iterative method that converges most of the time.

(1) Initialize $\mathcal{C}$ as the set of all strains in the community (including the new invader) and initialize the extinct set $\mathcal{E}$ as an empty set. V^* is the submatrix of interactions for the extant strains, and therefore contains the rows and columns corresponding to strains in $\mathcal{C}$.

(2) Invert V^* to find a fixed point of the dynamics via $\nu_i^* = \Upsilon^* \sum_j V^{*-1}_{ij}$ and $\sum_i \nu_i^* = 1$ which fixes Υ^*. In general some of the ν_i^* will be negative. Remove the set $\{i|\nu_i^* < 0\}$ from $\mathcal{C}$ and add it to $\mathcal{E}$.

(3) Calculate the invasion fitnesses ξ_j of each strain in $\mathcal{E}$ as $\xi_j = \sum_{k\in\mathcal{C}} V_{jk}\nu_k^* - \Upsilon^*$. Transfer the set of strains that could invade, $\{j|\xi_j > 0\}$, from $\mathcal{E}$ to $\mathcal{C}$.

(4) Iterate steps 2 and 3 until convergence to sets $\mathcal{C}$ and $\mathcal{E}$ such that $\nu_i^* > 0$ for all $i \in \mathcal{C}$ and $\xi_j < 0$ for all $j \in \mathcal{E}$.

One can then check stability of the final fixed point by computing the Jacobian of the dynamical equation. The vast majority of the time, the fixed point found in this way is stable.

When the algorithm does not converge within 10 iterations, we stop adding extinct strains with positive invasion fitnesses back into the community, and thereby converge to a fixed point with positive abundances, which is *invadable* by some of the strains that were declared extinct. This adjustment is most often deployed when $\mathcal{E}$ and $\mathcal{C}$ have similar sizes at the fixed point, which can happen for an initial assembled community with large S, or for the evolved state with small κ, where L is small and each invader can cause $O(L)$ extinctions. In these cases the iterative algorithm does not converge reliably. However, for small perturbations to the community, as in the case of primary interest, $\mathcal{E}$ is much smaller than $\mathcal{C}$ and the algorithm converges rapidly. Results from the error-prone regimes of evolution do not play a substantial part in any of our conclusions.

Note that this algorithm assumes invertibility of V^*. For the fully asymmetric model at steady state in the $Q \to \infty$ limit, $\overline{\Upsilon} = 0$ and therefore V^* has a zero eigenvalue: however for our numerics at finite Q, deviations in Υ allow this algorithm to be effective.

In the CR model, one cannot simply invert a matrix to find the fixed point for a subset of strains due to the nonlinearity in the interactions. Therefore, we use a numerical solver to find the fixed point for a given $\mathcal{C}$—the rest of the algorithm is the same as in the linearized and LV cases.

APPENDIX E: LINEARIZED MODEL ACROSS A RANGE OF PARAMETERS

To investigate the Red Queen phase in the absence of a Lyapunov function, we study the long-time averages $\overline{L}$, $\overline{\Upsilon}$, and $\overline{\log p_{\mathrm{inv}}}$ for ranges of the parameters κ, Σ, and mutant-parent correlation ρ. (Recall that we study $\overline{\log p_{\mathrm{inv}}}$ to focus on the *typical* p_{inv}.) Figures 10–12 display these quantities across a range of κ, Σ, and ρ.

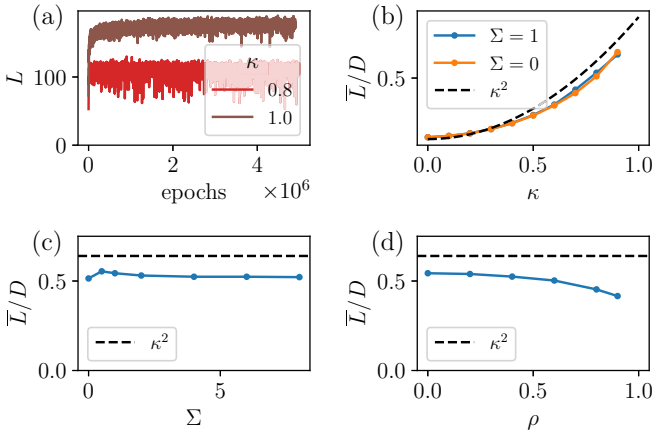


FIG. 10. Dynamics and steady-state of diversity in linearized model as a function of κ, Σ, and ρ. Default values $D = 200$, $\kappa = 0.8$, parent-mutant correlation $\rho = 0$, width of distribution of selective differences $\Sigma = 1$, with each varied as specified per panel. (a) Dynamics of L over evolutionary time for $\kappa = 0.8, 1$. Although L appears to be reaching steady state for both parameters, for $\kappa = 1$, Υ, and p_{inv} indicate it is not (Figs. 11 and 12). (b) Steady-state diversity. $\overline{L}$ normalized by its maximum, D, as a function of κ, for $\Sigma = 0$ and 1. Dotted line shows maximum diversity in *top-down assembled* community; in Red Queen phase, diversity is lower. (c) $\overline{L}$ depends only weakly on Σ. (d) L depends weakly on ρ for ρ not too near unity.

Like the CR model, the linearized model shows no Red Queen phase for perfectly symmetric interactions, but displays a Red Queen phase away from this special parameter choice. In Fig. 10(a) we plot the dependence of L on evolutionary time for two values of κ: 0.8 and 1. Starting from a single strain, there is a transient during which the diversity builds up, after which both trajectories flatten out. Figure 11(a) shows the trajectories of the Lagrange multiplier, Υ—analogous to the mean general fitness $\langle U \rangle$ in the CR model. For $\kappa = 0.8$, saturation of Υ indicates a Red Queen steady state and results in saturation of p_{inv} [Fig. 12(a)]. By contrast, for $\kappa = 1$, Υ continues to increase and p_{inv} concomitantly decreases: there is no Red Queen phase.

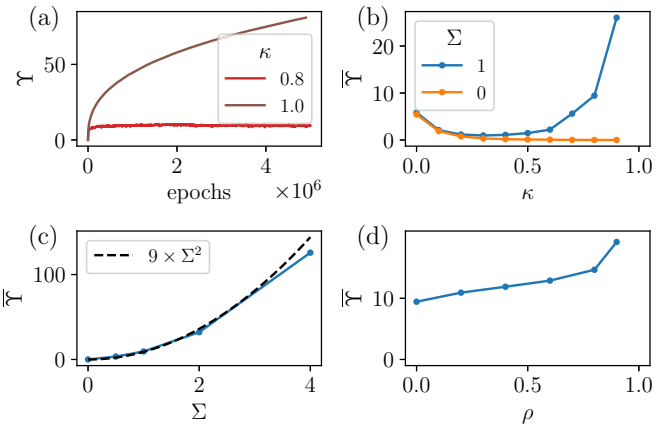


FIG. 11. Dynamics and steady state of Lagrange multiplier Υ—loosely community fitness—as a function of parameters for linearized resource model with default values as in Fig. 10. (a) Dynamics: for $\kappa = 1$, Υ (closely related to the Lyapunov function $\mathcal{L} = \Upsilon + \langle s \rangle$) is ever-increasing. For $\kappa = 0.8$, Υ reaches steady state $\overline{\Upsilon}$. (b) $\overline{\Upsilon}$ as a function of κ for $\Sigma = 0$ and 1. Only for $\Sigma > 0$ does the phenotype push into the tail of its distribution indicated by large Υ. Simulations are less reliable for small κ because L is smaller (see Appendix D). For $\kappa = 0.9$, $\Sigma = 1$, steady state not yet reached by end of simulation. (c) Dependence on Σ. Dotted line is rough fit to predicted form, $\overline{\Upsilon} \sim \Sigma^2$. At $\Sigma = 4$, Υ has not yet saturated. (d) Weak dependence of $\overline{\Upsilon}$ on ρ.

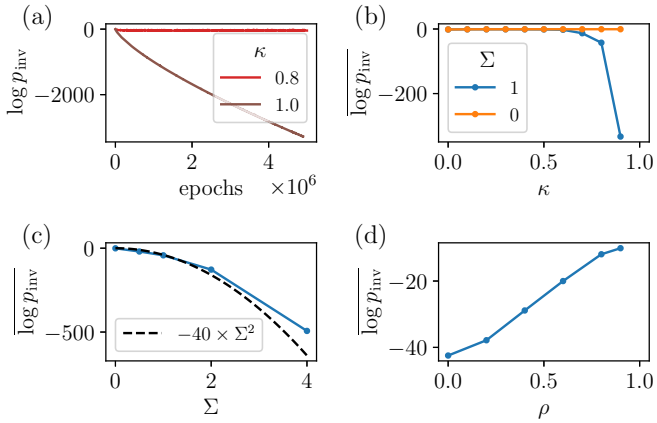


FIG. 12. Dynamics of invasion probability, and mean logarithm of invasion probability in steady state, $\overline{\log p_{\text{inv}}}$, as a function of parameters with default values as in Fig. 10. (a) Dynamics: p_{inv} saturates for $\kappa = 0.8$ but continues to decrease for $\kappa = 1$. (b) Dependence on κ. (c) Dashed line is rough fit to predicted $-\overline{\log p_{\text{inv}}} \sim \Sigma^2$, by analogy with assembled community [Eq. (5b)]. Simulations for $\Sigma = 4$ have not reached steady state and $\overline{\log p_{\text{inv}}}$ is overestimated (see Appendix N). (d) Invasion probability increases with ρ for ρ far from 1.

In Fig. 10(a), the observed $\overline{L}$ is compared with the curve $\kappa^2 D$ which is the maximum L in a top-down assembled community with F and G drawn from the Gaussian ensemble for the linearized model (see Sec. V A). We see that $\overline{L}$ from simulations is less than this naive upper bound, as will be discussed in Sec. V B. In Sec. V D we give scaling arguments for how $\overline{\Upsilon}$ and $\overline{\log p_{\text{inv}}}$ at steady state depend on κ and Σ: some of these scalings are compared with simulations in Figs. 11 and 12. Note that for the largest Σ the numerics have not reached steady state over the simulation time (Appendix N): this accounts for the difference between the numerical data and the theory.

APPENDIX F: LOTKA-VOLTERRA MODEL

In addition to the linearized model of Eq. (4), we have studied the generalized Lotka-Volterra (LV) model (or replicator equations) [62] which take the form

$$\frac{dv_i}{dt} = v_i\left(s_i + \sum_{j=1}^{S} V_{ij} v_j - \Upsilon(t)\right), \tag{F1}$$

with $\Upsilon(t) = \sum_i v_i(s_i + \sum_j V_{ij} v_j)$. This full-rank LV model has similar structure to the linearized model of Sec. IV A; the difference lies in the statistics and parametrization of the interaction matrix V. Instead of the $\mathcal{V} = -GF^{\intercal}$ with random

G and F of the linearized resource model, in this LV model the V_{ij} are Gaussian random variables with $\mathrm{E}[V_{ij}] = -Q\delta_{ij}$, $\mathrm{Var}[V_{ij}] = 1 + \gamma\delta_{ij}$, and $\mathrm{Cov}[V_{ij}, V_{ji}] = \gamma + \delta_{ij}$, with other covariances zero. Here $\gamma \in [-1, 1]$ is the symmetry parameter tuning between predator-prey-like interactions for $\gamma < 0$ and competitive interactions for $\gamma > 0$: a rough correspondence with the linearized model for positive γ is $\gamma \sim \kappa^2$. The negative diagonal term $-Q$ is a niche parameter that measures the strength of self-interaction relative to interstrain interaction, playing a similar role to $-\kappa\sqrt{D}$ in the linearized model, as noted in Sec. IV A. Our parametrization of the LV equations differs from a common parametrization in which interactions have nonzero mean [20]: here the V matrix captures *variations* in interaction magnitudes, with Υ taking the roles of the mean interaction strength and mean growth rate [24]. As in the linearized model, when the interactions are symmetric ($\gamma = 1$), $\mathcal{L} = \Upsilon + \langle s \rangle$ is a Lyapunov function of the dynamics, although in contrast with the linearized model [37], it can become a nonconvex function of the $\{\nu_i\}$ for large enough S [39].

The LV model has a similar ecological phase diagram to the CR model and linearized models, with a globally attracting fixed point of the dynamics when S is sufficiently small. The maximum L at a stable uninvadable fixed point scales as Q^2 with a γ-dependent coefficient [21]. Note that the scale of the drive distribution is $O(1/\sqrt{L}) = O(1/Q)$ in the LV model, as opposed to $O(\sqrt{D/L}) = O(1)$ in the linearized model, since the overall scale of the entries in the LV model is smaller by a factor of $\sqrt{D}$ than in the linearized model.

In the LV model, unlike in the linearized and CR models, strains cannot be characterized by finite-dimensional phenotypes. Although one can attempt to define a phenotype of strain i by s_i along with its corresponding row and column of V, this phenotype depends on the rest of the community and thus changes—or grows in dimension—as the community turns over.

APPENDIX G: PROPERTIES OF INVADERS

Here we define the specific parametrization of invaders that we have used for simulations.

In the LV model, the lack of phenotypes makes attempts to parametrize correlations between parent and mutants *ad hoc* for general γ [45], so we confine our study to the case of independent invaders. An invader strain I is generated by drawing a new row and column of V from the original Gaussian ensemble with $\mathrm{Corr}[V_{Ii}, V_{jI}] = \gamma\delta_{ij}$, and by drawing general fitness s_I from a Gaussian distribution with mean 0 and variance Σ^2. Therefore, the invasion probability of an independent invader is $p_{\mathrm{inv}} = 1 - \Phi[\Upsilon/(\sum_i \nu_i^2 + \Sigma^2)^{1/2}]$, where Φ is the standard normal CDF.

In models with phenotypes, such as the CR or linearized model, there is a natural way to generate mutants through small changes to the phenotype of an extant parent strain. This procedure creates a mutant whose interactions with all other strains are similar to those of its parent. In our Gaussian ensemble, the probability of any parent generating a mutant that can invade is the same—therefore we can, in an unbiased way, simply select a parent with probability proportional to its abundance.

The parameter $\rho \in [0, 1)$ describes the correlation between parent P and mutant M. The mutant phenotype is drawn from an ensemble that preserves $\mathrm{E}[G_{M\alpha}G_{P\beta}] = \mathrm{E}[F_{M\alpha}F_{P\beta}] = \rho\delta_{\alpha\beta}$, $\mathrm{E}[G_{M\alpha}F_{M\beta}] = \kappa\delta_{\alpha\beta}$, and $\mathrm{Corr}[s_M, s_P] = \rho$. We fix the desired correlations by drawing the phenotype of M as

$$G_{M\alpha} = \rho G_{P\alpha} + \sqrt{1-\rho^2}Z_\alpha, \tag{G1a}$$

$$F_{M\alpha} = \rho F_{P\alpha} + \sqrt{1-\rho^2}(\kappa Z_\alpha + \sqrt{1-\kappa^2}Z'_\alpha), \tag{G1b}$$

$$s_M = \rho s_P + \Sigma\sqrt{1-\rho^2}Z'', \tag{G1c}$$

where the $\{Z_\alpha\}$, $\{Z'_\alpha\}$, and Z'' are uncorrelated standard normal random variables. This choice ensures that the elements of the interaction matrix $V = -GF^\intercal$ have correlations $\mathrm{Corr}[V_{Mi}, V_{Pj}] = \mathrm{Corr}[V_{iM}V_{jP}] = \rho\delta_{ij}$ for $i, j \notin \{M, P\}$. The invasion probability of a mutant drawn from this ensemble is given by

$$p_{\mathrm{inv}} = 1 - \Phi\left[\sqrt{\frac{1-\rho}{1+\rho}}\frac{\Upsilon}{\sqrt{|\mathbf{R}|^2 + \Sigma^2}}\right], \tag{G2}$$

where Φ is the standard normal CDF. In this paper we restrict study to large effect mutations, with $(1-\rho^2)^{1/2} = O(1)$. The behavior as $\rho \to 1$, corresponding to small effect mutations, is both interesting and biologically relevant: we will discuss it in depth in future work.

APPENDIX H: EVOLUTION IN THE LOTKA-VOLTERRA MODEL

Given the apparent similarities between the CR and linearized models, it is natural to ask how the eco-evolutionary phenomenology in the random Lotka-Volterra model resembles or differs from these: in particular whether the difference in the ensemble of the interaction matrix between the LV and linearized models matters.

1. Red Queen phase

During the evolution, invading strains are chosen independently from the extant community and have correlations γ between their row and column added to the matrix. For a range of γ we find that the LV model exhibits a Red Queen phase analogous to those of the CR and linearized model, with the product of Q^2 and a function of γ determining the average number of strains in the steady state; we have also checked that a Red Queen phase obtains for nonzero Σ. Similarly to in the linearized model, the mean diversity $\overline{L}$ in the Red Queen phase is less than the maximum $\hat{L} = Q^2/(1+\gamma)^2$ from a top-down assembled community [Fig. 13(a)].

The gap between $\hat{L}$ and $\overline{L}$ diverges as $\gamma \to -1$, reflecting the fact that the fragility Ξ is numerically larger for $\gamma \to -1$, and so it drives an average of one strain extinct per invasion for L farther from $\hat{L}$. For $\Sigma = 0$ (as in these simulations) $\overline{\Upsilon}$ changes sign at $\gamma = 0$ [see Fig. 13(b)], at which we have a quantitative understanding of the Red Queen phase (Sec. V C). As in the linearized model, we expect weak dependence of $\overline{L}$ on Σ, and $\overline{\Upsilon} \sim \Sigma^2$ for Gaussian $p(s)$ with large Σ.

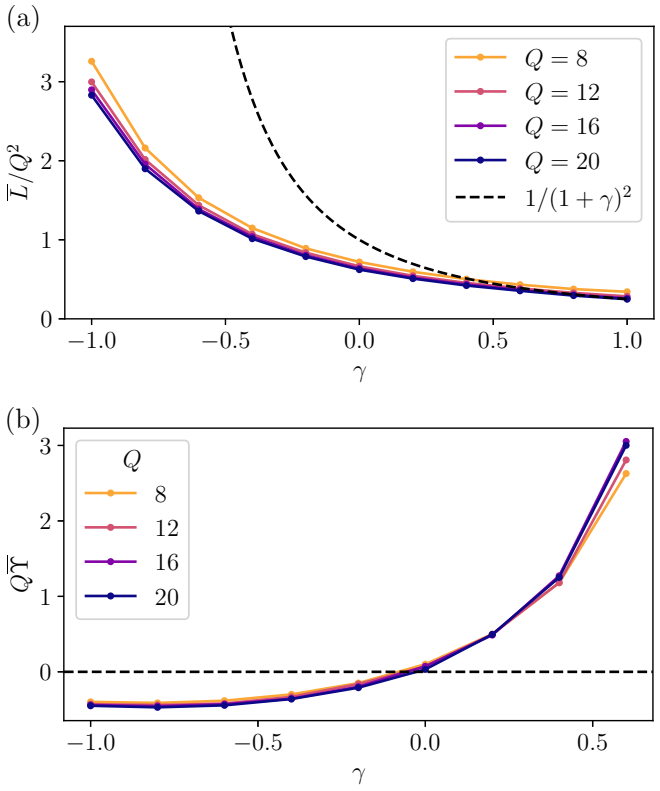


FIG. 13. Evolutionary dynamics of LV model (with $\Sigma = 0$) depend on γ, with Q setting an overall scale for both $\overline{L}$ and $\overline{\Upsilon}$. (a) Steady state $\overline{L}$ scaled by Q^2, in the LV model for $\Sigma = 0$ as a function of γ. Dashed line shows $\hat{L}/Q^2 = 1/(1+\gamma)^2$ from saturated assembled community [21]. When $\gamma > \gamma_c \approx 0.65$ the Red Queen phase breaks down, although L remains roughly constant, and its average is plotted. (b) $Q\overline{\Upsilon} = O(1)$ across a range of $\gamma < \gamma_c$ in the Red Queen phase. For $\gamma > \gamma_c$ there is no steady state, and $\overline{\Upsilon}$ is formally infinite. Dashed line shows $\hat{\Upsilon} = 0$ from saturated assembled community.

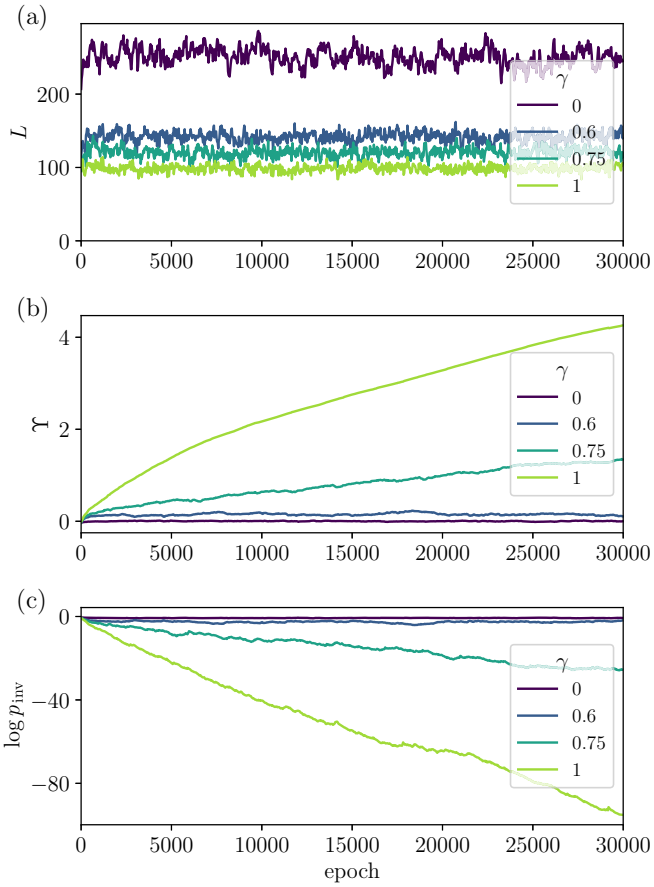


FIG. 14. Evolution in LV model ($Q = 20$, $\Sigma = 0$) undergoes a transition from Red Queen for $\gamma < \gamma_c \cong 0.65$ to oligarch phase for $\gamma > \gamma_c$. (a) L is roughly constant in both phases. (b) Υ fluctuates around the steady state in Red Queen phase but increases without bound in oligarch phase. (c) p_{inv} achieves a steady state in Red Queen phase but is ever-decreasing in oligarch phase.

2. Breakdown of Red Queen phase and oligarch phase

Although we find a robust Red Queen phase for a range of γ and Σ in the LV model, closer investigation of the effect of γ reveals a discontinuous transition around $\gamma_c \approx 0.65$ for $\Sigma = 0$. For $\gamma < \gamma_c$ a Red Queen phase obtains, while for $\gamma > \gamma_c$ we find the oligarch phase, in which the rate of successful invasion is ever-slowing and a few strains rise to high abundance.

Figures 7(a)–7(d) shows the trajectories of strain abundances in evolutionary time: here we show in Fig. 14 the corresponding trajectories of L, Υ, and p_{inv}. The oligarch phase is not a steady state and Υ runs away to arbitrarily large values, but L remains roughly constant, and it is this L that is plotted as $\overline{L}$ for $\gamma > \gamma_c$ in Fig. 13(a).

The rate of increase of Υ jumps discontinuously up from zero through the transition. This discontinuity suggests a first-order transition, motivating our investigation of hysteresis around the critical γ. One can start the evolution with $\gamma > \gamma_c$, and, after some period of time, reduce γ to below γ_c for the newly invading strains. Figure 7(e) shows strain abundance dynamics during such a "quench" from $\gamma = 1$ to $\gamma = 0.5$, and Fig. 15 shows the dynamics of L, Υ, and p_{inv} through this process, compared with the behaviors for constant $\gamma = 0.5$ and $\gamma = 1$. If one looks only at L, the hysteresis is not apparent, with the postquench L similar to L for a community that has always had $\gamma = 0.5$. Is it possible that, as for $\gamma = 0$, the fragility only depends on L but not the abundance distribution? If so, the fragility could set L in steady state with little dependence on the history.

In contrast with L, Υ, and p_{inv} show clearly that the postquench steady state is significantly different from a Red Queen state that has been evolving with $\gamma = 0.5$ from its beginning. This postquench steady state displays properties of both the oligarch phase and the Red Queen phase, and may therefore be a useful route to better understand the oligarch phase.

APPENDIX I: CAVITY CALCULATION FOR ASSEMBLED COMMUNITY

Here we carry out the static cavity calculation for a top-down assembled community in the linearized model: our results and approach are similar to that of Ref. [22]. An analogous calculation for the LV model proceeds in a similar

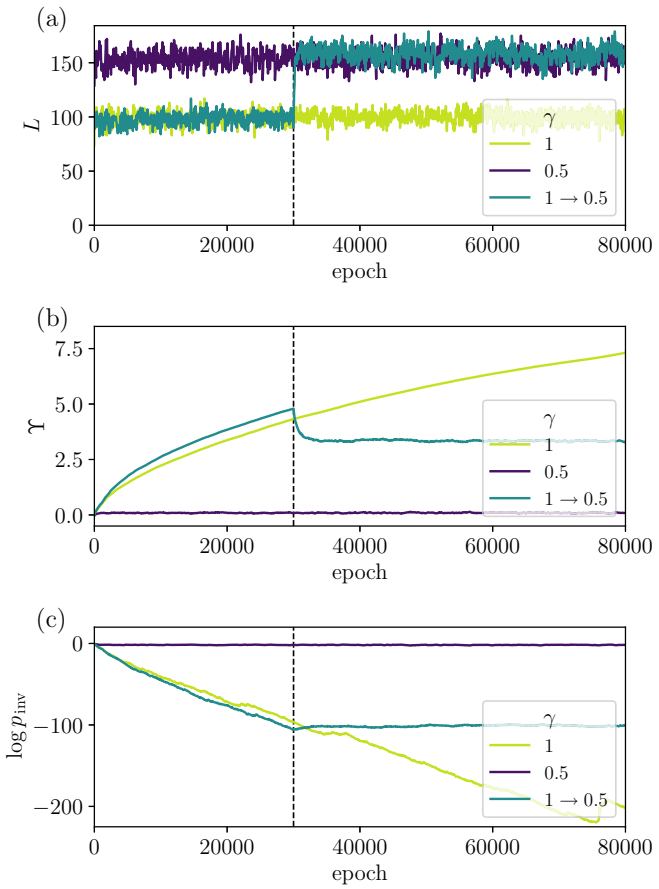


FIG. 15. Transition between Red Queen and oligarch phases is first order and exhibits hysteresis and phase coexistence when the system is tuned across the transition, shown for $Q = 20$, $\Sigma = 0$. γ is reduced from 1 to 0.5 after 3×10^4 epochs of evolution. (a) The trajectory of L over such an evolution, shown alongside L for constant $\gamma = 0.5$ and constant $\gamma = 1$. The value of L is indistinguishable in the postquench phase (after a short transient) from the $\gamma = 0.5$ Red Queen phase. (b), (c) Dynamics of Υ and p_{inv} show a strong departure after the quench from their values when $\gamma = 0.5$ throughout.

manner [20,21,24], and we summarize it at the end of this section. We consider a large community of strains at its fixed point, and add in another strain and resource, treating these as a small perturbation and then enforcing self-consistency between the statistics of this newly added strain and resource and the rest of the community. We use the convention that repeated indices are summed over, unless otherwise stated.

1. Linearized model

The strains obey the dynamical equation

$$\frac{\dot{\nu}_i}{\nu_i} = s_i + \mathcal{V}_{ij}\nu_j - \Upsilon \tag{I1}$$

for $j \in \{1,\dots,S\}$, where $\mathcal{V}_{ij} = -G_{i\alpha}F_{j\alpha}$ for $\alpha \in \{1,\dots,D\}$ and $\mathrm{E}[G_{i\alpha}F_{j\beta}] = \kappa\delta_{ij}\delta_{\alpha\beta}$ with $\Upsilon = \nu_i(s_i + \mathcal{V}_{ij}\nu_j)$. Defining resources $R_\alpha = -F_{j\alpha}\nu_j$, the fixed-point abundances satisfy

$$0 = s_i + G_{i\alpha}\tilde{R}_\alpha - \Upsilon, \quad \tilde{R}_\alpha = -F_{j\alpha}\tilde{\nu}_j, \tag{I2}$$

where tildes on dynamical variables ν and R indicate the fixed-point values prior to adding a new strain and resource. Now we add in a new strain and resource with fixed point abundances ν_0 and R_0, respectively. Defining quantities $\zeta_0 = G_{0\alpha}\tilde{R}_\alpha$ and $\phi_0 = -F_{j0}\tilde{\nu}_j$, we have

$$0 = s_0 + \zeta_0 + G_{0\alpha}\delta R_\alpha - \Upsilon, \quad R_0 = \phi_0 - F_{j0}\delta\nu_j. \tag{I3}$$

If we define "local susceptibilities" $\eta_\alpha = dR_\alpha/d\phi_\alpha$ and $\chi_i = d\nu_i/d\zeta_i$, then the changes in R_α and ν_j are $\delta R_\alpha = -\eta_\alpha F_{0\alpha}\nu_0$ and $\delta\nu_j = \chi_j G_{j0}R_0$. (Here repeated indices are not summed.)

Assuming self-averaging (distributions concentrating around their mean for large systems), the new strain and resource abundances are

$$\nu_0 = \frac{\zeta_0 + s_0 - \Upsilon}{\kappa H}, \quad R_0 = \frac{\phi_0}{1+\kappa X}, \tag{I4}$$

where $H = \sum_\alpha \eta_\alpha$ and $X = \sum_i \chi_i$ are "global" susceptibilities.

Since ζ and ϕ are weighted sums of the G and F matrix elements respectively, they are Gaussian distributed with mean 0. We denote their variances as σ_ζ^2 and σ_ϕ^2, respectively. Let us take the s to be Gaussian distributed with mean zero and variance Σ^2, and define $\sigma_\xi^2 = \sigma_\zeta^2 + \Sigma^2$ and $Y = \Upsilon/\sigma_\xi$. We can now write self-consistency conditions for the variables σ_ϕ, σ_ξ, H, X, and Y, which require that the distributions of the added strain and resource abundances are the same as the distributions in the original community: in other words, Eq. (I4) is true for all strains, not just the cavity strain. Defining partial moments

$$\mathcal{W}_m(Y) = \frac{1}{\sqrt{2\pi}}\int_0^\infty x^m e^{-(x+Y)^2/2}dx, \tag{I5}$$

our self-consistency conditions are

$$X = \frac{S}{\kappa H}\mathcal{W}_0(Y), \tag{I6a}$$

$$1 = \frac{S\sigma_\xi}{\kappa H}\mathcal{W}_1(Y), \tag{I6b}$$

$$\sigma_\phi^2 = \frac{S\sigma_\xi^2}{\kappa^2 H^2}\mathcal{W}_2(Y), \tag{I6c}$$

$$H = \frac{D}{1+\kappa X}, \tag{I6d}$$

$$\sigma_\xi^2 = \frac{D\sigma_\phi^2}{(1+\kappa X)^2} + \Sigma^2. \tag{I6e}$$

These equations can be reduced to

$$X = \frac{S\mathcal{W}_0(Y)}{\kappa[D - S\mathcal{W}_0(Y)]}, \tag{I7a}$$

$$H = D - S\mathcal{W}_0(Y), \tag{I7b}$$

$$\sigma_\xi = \kappa\frac{D - S\mathcal{W}_0(Y)}{S\mathcal{W}_1(Y)}, \tag{I7c}$$

along with an equation for Y alone:

$$\frac{S\mathcal{W}_2(Y)}{\kappa^2 D} = 1 - \left[\frac{\Sigma S\mathcal{W}_1(Y)}{\kappa[D - S\mathcal{W}_0(Y)]}\right]^2. \tag{I8}$$

A solution to these equations gives us the distributions of the strain and resource abundances in the community, as well as

the susceptibilities H and X, and the scaled Lagrange multiplier Y. This mean-field solution is expected to describe the unique stable uninvadable fixed point as long as it is stable to a small random perturbation on each of the strains' growth rates.

2. Fragility and instability

To look at the nonlinear instability that signals the breakdown of the mean-field theory, we add a random field h_i to each of the biases and look at the change $\delta\xi_0$ in the bias of strain zero. We find that

$$\delta\xi_0 = \eta_\alpha G_{0\alpha} F_{j\alpha}\chi_j \delta\xi_j + h_0, \tag{I9}$$

where the sum is over both α and j. Then in evaluating the mean square of both sides of this equation we have

$$\mathrm{E}\left[\left(\eta_\alpha G_{0\alpha} F_{j\alpha}\chi_j\delta\xi_j\right)^2\right] = \sum_\alpha \eta_\alpha^2 \sum_j \chi_j^2\,\mathrm{E}\left[\delta\xi_i^2\right]. \tag{I10}$$

Assuming that the strains are statistically identical, we can drop the strain indices to get

$$\mathrm{E}[\delta\xi^2] = \frac{\mathrm{E}[h^2]}{1-\sum_\alpha \eta_\alpha^2 \sum_j \chi_j^2} \tag{I11a}$$

$$\Rightarrow \sum_i \mathrm{E}\left[\delta\nu_i^2\right] = \frac{\sum_j \chi_j^2\,\mathrm{E}[h^2]}{1-\sum_\alpha \eta_\alpha^2 \sum_j \chi_j^2} = \Xi \times \mathrm{E}[h^2], \tag{I11b}$$

where we have defined the *fragility* Ξ as the factor relating the variance of the random fields to the sum of the squared changes in abundances that these produce. Since $\eta_\alpha = H/D$ and $\chi_j = X/S\mathcal{W}_0(Y)$ in the mean-field solution, we can write the fragility in terms of $L = S\mathcal{W}_0(Y)$ as

$$\Xi = \frac{L/\kappa^2 H^2}{1 - L/\kappa^2 D} = \frac{LD}{(D-L)^2(\kappa^2 D - L)}. \tag{I12}$$

An instability signaling a phase transition occurs when the fragility diverges, at $L = \kappa^2 D$. We could have (more naturally) defined the fragility as the response of the strain abundances to adding a random field to the resource fields $\{\phi_\alpha\}$, but the form of divergence at the phase transition is the same in both cases (and in LV model, only the fragility we use here is defined). The location of the divergence implies that the number of surviving strains is $\hat{L} = \kappa^2 D$ at the transition. Substituting this relation into Eq. (I8) tells us that $\hat{Y}$ is the solution to

$$\frac{\mathcal{W}_2(\hat{Y})}{\mathcal{W}_0(\hat{Y})} = 1 - \left[\frac{\mathcal{W}_1(\hat{Y})\Sigma\kappa}{\mathcal{W}_0(\hat{Y})(1-\kappa^2)}\right]^2. \tag{I13}$$

For $\Sigma = 0$, the solution to this equation is $\hat{Y} = 0$ which implies $\hat{S} = 2\kappa^2 D$. For $\Sigma \gtrsim 1$, we have $\hat{Y} \approx \frac{\kappa\Sigma}{1-\kappa^2}$ and $\hat{\sigma}_\xi \approx \Sigma$, meaning that $\hat{\Upsilon} \approx \frac{\kappa\Sigma^2}{1-\kappa^2}$ and $\hat{S} \approx \sqrt{2\pi}\kappa^2 D\hat{Y} e^{\hat{Y}^2/2} \sim D\exp[\frac{\kappa^2\Sigma^2}{2(1-\kappa^2)^2}]$. Note that $\hat{L}/\hat{S}$, the probability of a particular initial strain surviving in the assembled community, gives the probability that an independent invader could invade the saturated community. The scaling of $\hat{\Upsilon}$ and $\log\hat{S}$ with Σ and κ is reflected in the scaling forms of $\overline{\Upsilon}$ and $\overline{\log p_{\mathrm{inv}}}$ in the Red Queen steady state.

Note that the linear susceptibility X diverges when $L = D$. The nonlinear instability occurs for lower S than the linear instability for $\kappa < 1$, meaning that the cavity solution becomes invalid before it displays a linear instability. The nonlinear and linear instabilities occur in the same place for $\kappa = 1$, marking the onset of a shielded phase when $\Sigma = 0$ [37]. With $\Sigma > 0$, $\hat{S}$ diverges for $\kappa \to 1$, indicating the marginality of the shielded phase which occurs only for $\Sigma = 0$, $\kappa = 1$.

3. Lotka-Volterra model

We now summarize the cavity calculation for the LV model. In this case the cavity method yields, for the fixed-point abundance of the cavity strain,

$$\nu_0 = \frac{\zeta_0 + s_0 - \Upsilon}{Q - \gamma X}, \tag{I14}$$

where $\zeta_0 = \sum_j V_{0j}\tilde{\nu}_j$, and $X = \sum_i \chi_i = \sum_i d\nu_i/d\zeta_i$. Note the similarity with the expression from the linearized model

$$\nu_0 = \frac{\zeta_0 + s_0 - \Upsilon}{\kappa D - \kappa^2 H X} \tag{I15}$$

[a consequence of Eqs. (I4) and (I7)] if one identifies $\kappa\sqrt{D}$ with Q, κ^2 with γ, and HX with $\sqrt{D}X$, which account for the integrated-out resources as well as the factor of $\sqrt{D}$ between the LV and linearized models.

The self-consistency equations for the LV model with Gaussian distributed s_i of variance Σ^2, for unknowns σ_ξ, Y, and X, are then

$$X = \frac{S}{Q-\gamma X}\mathcal{W}_0(Y), \tag{I16a}$$

$$1 = \frac{S\sigma_\xi}{Q-\gamma X}\mathcal{W}_1(Y), \tag{I16b}$$

$$\sigma_\xi^2 = \frac{S\sigma_\xi^2}{(Q-\gamma X)^2}\mathcal{W}_2(Y) + \Sigma^2, \tag{I16c}$$

where, as before, $Y = \Upsilon/\sigma_\xi$. These equations can be reduced to

$$X = \frac{Q - \sqrt{Q^2 - 4\gamma S\mathcal{W}_0(Y)}}{2\gamma}, \tag{I17}$$

$$\sigma_\xi = \frac{Q + \sqrt{Q^2 - 4\gamma S\mathcal{W}_0(Y)}}{2S\mathcal{W}_1(Y)}, \tag{I18}$$

along with an equation for Y alone:

$$\left[\frac{Q + \sqrt{Q^2 - 4\gamma S\mathcal{W}_0(Y)}}{2}\right]^2 = S\mathcal{W}_2(Y) + [\Sigma S\mathcal{W}_1(Y)]^2. \tag{I19}$$

The fragility can be calculated from local susceptibility $\chi_i = 1/(Q-\gamma X)$ in terms of $L = S\mathcal{W}_0(Y)$:

$$\Xi = \frac{L}{(Q-\gamma X)^2 - L}, \tag{I20}$$

implying a divergence of Ξ at $\hat{L} = Q^2/(1+\gamma)^2$.

We note that for the fully asymmetric LV model, the fragility can be calculated directly due to the simple relationship between ν_i and ξ_i. For each strain with positive bias, $\nu_i = \xi_i/Q$ so the susceptibility to a change in growth rate is $1/Q$ (Appendix I). However, the change in each ν_i causes a

small random change of the biases of all the other strains, which increases the effective variance of the random field, and so on. Summing up these effects [45] yields a total effect of random h_j given by $\sum_i(\delta\nu_i)^2 = \Xi\,\mathrm{Var}[h_j]$ with the fragility $\Xi = \frac{L/Q^2}{1-L/Q^2} = \frac{L}{Q^2-L}$. This result does not depend on how the community is formed, or on the selective differences if they are present. It shows right away the maximum size of a stable community is $\hat{L} = Q^2$, which is saturated for an assembled community. For the evolved community, Ξ cannot be large—otherwise multiple strains would be driven extinct by each invader. Thus $\overline{L} < \hat{L}$, as we find in numerics.

APPENDIX J: EXPONENTIAL DISTRIBUTION OF GENERAL FITNESSES

To analyze the transient behavior of L as $\Delta(\Upsilon)$ (the local scale of the s_i distribution) decreases during the evolutionary transient, it is useful to investigate the behavior of the assembled community with an exponential distribution of positive s_i given by $p(s) = \frac{1}{\Delta}e^{-s/\Delta}\Theta(s)$. We define the variable $x_i \equiv s_i + \zeta_i$. The self-consistency conditions then deal with the distribution of x_i, which is

$$p_x(x) = \frac{1}{\Delta}\exp\left[-\frac{x}{\Delta} + \frac{\sigma_\zeta^2}{2\Delta^2}\right]\Phi\left(\frac{x}{\sigma_\zeta} - \frac{\sigma_\zeta}{\Delta}\right), \tag{J1}$$

where Φ is the standard normal CDF and σ_ζ^2 is the variance of the drives in the *assembled* community. If we define partial moments

$$\tilde{\mathcal{W}}_m(\Upsilon, \sigma_\zeta) = \int_0^\infty y^m p_x(y+\Upsilon)dy, \tag{J2}$$

then the self-consistency conditions can be written down in a similar fashion to those for the Gaussian $p(s)$ in Eq. (I6). For the LV model, these conditions are

$$X = \frac{S}{Q-\gamma X}\tilde{\mathcal{W}}_0(\Upsilon, \sigma_\zeta), \tag{J3}$$

$$1 = \frac{S}{Q-\gamma X}\tilde{\mathcal{W}}_1(\Upsilon, \sigma_\zeta), \tag{J4}$$

$$\sigma_\zeta^2 = \frac{S}{(Q-\gamma X)^2}\tilde{\mathcal{W}}_2(\Upsilon, \sigma_\zeta). \tag{J5}$$

For large Υ, the lower limit of the exponential $p(s)$ does not matter and the distribution $p_x(x)$ is roughly an unbounded exponential. Therefore, we have

$$\tilde{\mathcal{W}}_m(\Upsilon, \sigma_\zeta) \approx m!\Delta^m \exp\left[-\frac{\Upsilon}{\Delta} + \frac{\sigma_\zeta^2}{2\Delta^2}\right], \tag{J6}$$

which allows us to write the self-consistency conditions just in terms of the number of extant strains $L = \tilde{\mathcal{W}}_0(\Upsilon, \sigma_\zeta)$, regardless of S and the normalization of $p(s)$:

$$X = \frac{L}{Q-\gamma X}, \tag{J7}$$

$$1 = \frac{\Delta L}{Q-\gamma X}, \tag{J8}$$

$$\sigma_\zeta^2 = \frac{2\Delta^2 L}{(Q-\gamma X)^2}. \tag{J9}$$

In the large-Δ regime, we therefore have $X \approx 1/\Delta$, implying that, to leading order in Δ, the diversity is given by $L(\Delta) \approx Q/\Delta$. A similar analysis in the linearized model yields $L(\Delta) \approx \kappa D/\Delta$.

As discussed in Sec. V D, the analysis of the exponential $p(s)$ also describes the steady-state behavior for $p(s)$ decaying faster than exponentially, and Σ large. In that case, Δ is a function of Υ and the distribution of s_i has pushed far enough into the tail of $p(s)$ to look locally exponential around Υ.

Note that $\hat{\Upsilon}$ and $\hat{S}$ diverge for sufficiently large $\Delta > \Delta_c$ whether the critical Δ, denoted Δ_c, is proportional to—with an $O(1)$ factor—$(1+\gamma)/Q$ in the LV model and $1/\kappa$ in the linearized model. This is because, as Δ increases, the maximum L—given by $X(Q-\gamma X)$ and $X\kappa D/(1+\kappa X)$ in the LV and linearized models, respectively—becomes smaller than $Q^2/(1+\gamma)^2$ or $\kappa^2 D$, respectively, and so Ξ cannot diverge. Therefore, for an exponential $p(s)$, sufficiently large Δ eliminates the instability of an assembled community.

APPENDIX K: DYNAMICAL CAVITY ANALYSIS FOR FULLY ASYMMETRIC LOTKA-VOLTERRA MODEL

Here we carry out the dynamical cavity analysis for the evolutionary steady state of the LV model with $\gamma = 0$, $\Sigma = 0$. In particular we calculate the temporal correlation function of the cavity strain in Eq. (7) and derive self-consistency conditions to find the statistical properties of the evolution of strain abundances. Throughout, we suppress the index 0 of the cavity strain.

1. Solution for correlation function

Because of the absorbing boundary condition at $\nu = 0$, we can use the solution of the Ornstein-Uhlenbeck process [Eq. (8)] together with the method of images to write down the distribution of $\nu(T)$ given that it started at some initial value $\nu(\mathbb{T})$ at invasion time $T = \mathbb{T}$. For simplicity we set $\mathbb{T} = 0$: integrating the correlation function over $\mathbb{T}$—as we later do—is the same as integrating over T. We use the fact that, for $\gamma = \Sigma = \overline{\Upsilon} = 0$, we have $\nu = \zeta/Q$. Therefore, we can write the distribution of $\nu(T)$ conditioned on $\nu(0)$ as

$$p^*(\nu(T)|\nu(0)) = \frac{1}{\sqrt{2\pi}\sigma(T)}\left[\exp\left(-Q^2\frac{[\nu(T)-\nu(0)e^{-bT}]^2}{2\sigma(T)^2}\right) - \exp\left(-Q^2\frac{[\nu(T)+\nu(0)e^{-bT}]^2}{2\sigma(T)^2}\right)\right]$$

$$\text{with}\quad \sigma^2(T) = \frac{J}{b}(1-e^{-2bT}), \tag{K1}$$

where we have used the notation p^* to denote the probability density with an absorbing boundary condition at $\nu = 0$: the integral of p^* over $\nu(T)$ is the survival probability of the strain as a function of time T. We calculate the expectation $\langle\nu(T)\nu(T+\tau)\rangle$ over the joint distribution $p^*(\nu(T), \nu(T+\tau)|\nu(0)) = p^*(\nu(T)|\nu(0)) \times p^*(\nu(T+\tau)|\nu(T))$. We then integrate over all times $T \in [0, \infty)$ and over all $\nu(0) \in [0, \infty)$ distributed according to

$$p(\nu(0)) = \frac{r}{\sqrt{2\pi}\sigma_\zeta}\exp\left[-\frac{Q^2\nu(0)^2}{2\sigma_\zeta^2}\right]\Theta[\nu(0)], \tag{K2}$$

where $\sigma_\zeta = \sqrt{J/b}$ is the standard deviation of the incoming drive distribution, and r is the rate of attempted invasions. $r = 2$ in our case since time is in units of successful invasions, and $\overline{\Upsilon} = 0$ so each invader has probability 1/2 of successfully invading. The integrals over T and $\nu(0)$ correspond to summing over all the strains in the community and can be done exactly, yielding

$$C(\tau) = \mathrm{E}\left[\langle \nu(T)\nu(T+\tau)\rangle\right]_{T,\nu(0)} \tag{K3a}$$

$$= \frac{r\sigma_\zeta^2}{4Q^2 b}(1 + 2\log 2)e^{-b\tau}. \tag{K3b}$$

We now assert that the dynamics of the cavity drive are the same as those of the drives in the community. Since the cavity drive can fluctuate in sign, its correlation function is $C(\tau) = \sigma_\zeta^2 e^{-b\tau}$, from an unconditioned Ornstein-Uhlenbeck process at stationarity. Equating our two expressions for $C(\tau)$ gives

$$1 = \frac{r}{4Q^2 b}(1 + 2\log 2). \tag{K4}$$

We can get another self-consistency condition by enforcing $\sum_j \nu_j = 1$. Note that $\langle \nu(T)|\nu(0)\rangle = \nu(0)e^{-bT}$, where the expectation is taken with an absorbing boundary condition at $\nu = 0$. Integrating this expression over the past and over the half-Gaussian distribution of starting positions (as before) gives $r\sigma_\zeta = Qb\sqrt{2\pi}$.

The number of strains in steady state, $\overline{L}$, is determined by the fact that the survival probability up to time T is $\langle \Theta[\nu(T)]|\nu(0)\rangle = \mathrm{erf}[Q\nu(0)e^{-bT}/\sqrt{2}\sigma(T)]$. Integrating over all past times and over the distribution of $\nu(0)$ gives $\overline{L} = \frac{r}{2b}\log 2$. Now we have two self-consistency conditions for b and σ_ζ, along with an equation for $\overline{L}$:

$$r\sigma_\zeta = Qb\sqrt{2\pi}, \tag{K5}$$

$$4Q^2\frac{b}{r} = 1 + 2\log 2, \tag{K6}$$

$$\overline{L} = \frac{r}{2b}\log 2. \tag{K7}$$

We can solve these equations in terms of Q. The most easily testable prediction from numerics is

$$\overline{L} = \frac{2\log 2}{1 + 2\log 2}Q^2 \cong 0.58Q^2. \tag{K8}$$

We also find that $\sigma_\zeta = \sqrt{2\pi}\lambda/Q$, and $b = \lambda r/Q^2$ where $\lambda = (1 + 2\log 2)/4$, implying that $J = 2\pi r\lambda^3/Q^4$. Therefore, the theory prediction for the correlation function of the cavity drive is

$$C(\tau) = \frac{2\pi\lambda^2}{Q^2}\exp\left[-\frac{\lambda}{Q^2}r\tau\right]. \tag{K9}$$

As stated previously, $r = 2$ since this makes the integral of the half-Gaussian incoming drive distribution (and therefore the rate of incoming strains) equal to one.

Figure 16 compares these analytical results to numerics on the $\gamma = 0$, $\Sigma = 0$ LV model. As Q increases in Fig. 16(a), the numerics get closer to the large-$\overline{L}$ regime in which the theory holds, and $\overline{L}/Q^2$ in the steady-state approaches $(2\log 2)/(1 + 2\log 2)$ from above. Figure 16(b) shows the scaled correlation function $Q^2C(\tau)$ versus scaled time τ/Q^2, which agrees well with Eq. (K9).

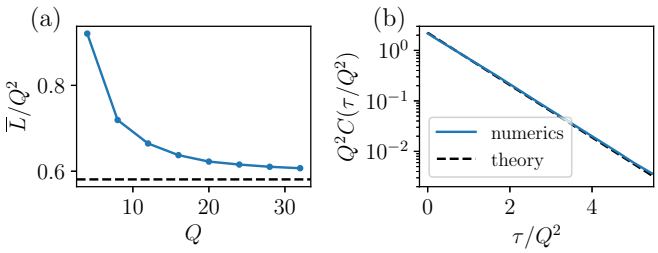


FIG. 16. Predictions of evolutionary cavity calculation and simulations for Red Queen state of $\gamma = 0$, $\Sigma = 0$, LV model with independent invaders. (a) Number of strains as a function of Q converges to predicted $\overline{L} \approx 0.58\,Q^2$ (dotted black line) as $Q \to \infty$ [Eq. (K8)]. (b) Temporal correlation function $C(\tau) = \sum_j \nu_j(T)\nu_j(T+\tau)$ agrees with Eq. (K9) with no fit parameters (here $Q = 20$).

2. Lifetime distribution

The distribution of strain lifetimes can be computed using known results from the theory of Ornstein-Uhlenbeck (OU) processes—which can also be derived from the method of images. A useful result [63] is that the distribution of survival times $\mathcal{T}$ for an OU process $\dot{x} = -bx + \sqrt{2J}\eta(T)$ starting at position $x_0 > 0$ and ending when it hits the origin is

$$p_{\mathcal{T}}(\mathcal{T}|x_0) = \frac{2b^{3/2}x_0}{\sqrt{2\pi J}}\frac{e^{2b\mathcal{T}}}{(e^{2b\mathcal{T}} - 1)^{3/2}}\exp\left[-\frac{bx_0^2}{2J(e^{2b\mathcal{T}} - 1)}\right]. \tag{K10}$$

We can then integrate over the distribution

$$p_{x_0}(x_0) = \frac{2\Theta(x_0)}{\sqrt{2\pi}\sigma_\zeta}e^{-x_0^2/2\sigma_\zeta^2}, \quad \text{where} \quad \sigma_\zeta = \sqrt{J/b}. \tag{K11}$$

We obtain $p_{\mathcal{T}}(\mathcal{T}) = 2b(e^{(2b\mathcal{T})} - 1)^{-1/2}/\pi$, as quoted in Eq. (11).

3. Abundance distribution

The steady-state distribution of abundances can also be computed from our self-consistent solution. The abundance distribution obeys a Fokker Planck equation (with the appropriate boundary conditions) that captures the dynamics of the drives from Eq. (8) and includes a *source* term for incoming invaders. Therefore, in the steady state, the condition for stationarity of the distribution $P(\nu)$ is

$$\frac{b\sigma_\zeta^2}{Q^2}P''(\nu) + b[P(\nu) + \nu P'(\nu)] + \frac{rQ}{\sqrt{2\pi}\sigma_\zeta}\exp\left(-\frac{\nu^2Q^2}{2\sigma_\zeta^2}\right) = 0, \tag{K12}$$

with boundary conditions $P(0) = P(\nu \to \infty) = 0$. Equation (K12) contains effects of the systematic and random parts in $\delta\nu$ as well as the input from invading strains.

If we change variables to $z = Q^2\nu$ and $\Pi(z) = \frac{1}{Q^4}P(\nu)$ so that $\overline{L} = Q^2\int dz\Pi(z)$, we find that the number density $\Pi(z)$ of the rescaled abundances z obeys the Q-independent differential equation

$$2\pi\lambda^3\Pi''(z) + \lambda[\Pi(z) + z\Pi'(z)] + \frac{1}{2\pi\lambda}\exp\left(-\frac{z^2}{4\pi\lambda^2}\right) = 0, \tag{K13}$$

with $\lambda = (1 + 2\log 2)/4$. The solution to this equation with boundary conditions $\Pi(0) = 0$, $\Pi(\infty) = 0$ is plotted in Fig. 5 as the theoretical prediction for the distribution of scaled abundances in the evolved community.

APPENDIX L: GENERAL CAVITY FORMULATION

We now generalize the cavity analysis away from $\gamma = 0$, $\Sigma = 0$. For $\gamma \neq 0$ in the LV model the analysis is particularly complicated. (For $\kappa > 0$ in the linearized model it is even more so as both resources and strain dynamics have to be analyzed as in the ecological cavity calculation of Appendix I.) In general, the bias and abundance of a focal strain are no longer simply related as they were in the $\gamma = 0$ case: due to the feedback of the strain on itself via its influence on the other strains, $\nu_0(T)$ will depend both on its bias $\xi_0(T)$ and an additive term arising from the feedback integrated over all times since its invasion. Therefore, the analog of Eq. (6) for $\gamma \neq 0$ is

$$\nu_0(T) = \frac{\phi_0(T)}{Q}\mathbb{1}[\phi_0(T') > 0\ \forall\ T' \in (\mathbb{T}_0, T)], \tag{L1}$$

where

$$\phi_0(T) = \xi_0(T) + \gamma\int_{\mathbb{T}_0}^{T} dT' M(T - T')\nu_0(T'), \tag{L2}$$

with $M(\tau)$ a memory kernel that must be determined self-consistently. It will include an instantaneous—on evolutionary timescales—part from the invader's *ecological* feedback on its own abundance, in addition to a slow memory part that extends over many epochs [18]. Both the memory effects and the $s_0 - \Upsilon$ part of ξ_0 mean that the cavity strain will not go extinct when ζ_0 goes through zero but instead at some other time when $\phi_0(T)$ first goes negative.

The statistics of the drives $\zeta_0(T)$ are determined by the extant community, but $C(\tau)$ is no longer a simple exponential. Nevertheless, one can write an effective stochastic equation for ζ_0:

$$\frac{d\zeta_0}{dT} = -\int_{\mathbb{T}_0}^{T} dT' B(T - T')\zeta_0(T') + \sqrt{2J}\eta_0(T), \tag{L3}$$

with $B(\tau)$ another memory kernel and $\eta_0(T)$ still white noise which represents the diffusion-like effects of an invader on the other strains discussed in Sec. V B. The Fourier transform of B is related, via a Weiner-Hopf factorization, to the Fourier transform of C: denoting the transforms by tildes, we have

$$\tilde{C}(\omega) = \frac{2J}{[-i\omega + \tilde{B}(\omega)][+i\omega + \tilde{B}(-\omega)]}, \tag{L4}$$

with $\tilde{B}(\omega)$ analytic in the upper half of the complex plane.

In principle, to proceed one assumes a $C(\tau)$—or equivalently a $B(\tau)$ and J—as well as Υ [and if $\gamma \neq 0$, also $M(\tau)$]. Then one computes $\langle\nu_0(T)\nu_0(T + \tau)\rangle$ conditioned on ν_0 invading at time $\mathbb{T}_0$. By averaging over $\nu_0(\mathbb{T}_0)$, $\mathbb{T}_0$, and s_0, one can enforce the self-consistency of $C(\tau)$ using Eq. (7), as for the $\gamma = 0$ case. Then $M(\tau)$ is self-consisted from calculating the response of ν_0 to a change in its bias, and Υ is self-consisted to ensure $\sum_i \nu_i = 1$. In the linearized model with resources, the resources have to be handled by a cavity method as well but their effects can be integrated out to give a more complicated relation between strain abundances and M, with Q replaced by $\kappa\sqrt{D}$ and γ replaced by κ^2 (Appendix I).

Unfortunately, even numerical analysis of the self-consistency conditions to determine the memory functions M and B is likely challenging [53]. However, the above structure allows us to glean some insights.

In the absence of selective differences, all strains are statistically similar and the systematic decrease in ζ_0 should make the strain lifetime distribution decay exponentially at long times. This suggests that both the memory kernels M and B will decay exponentially at long times, as will the correlation function of the drives, $C(\tau)$. To explore the effects of feedback, we carried out simulations of the $\Sigma = 0$ LV model for $\gamma = -0.8$, 0, and 0.5, and plotted $C(\tau)$ in Fig. 19. We see that $C(\tau)$ indeed decays exponentially at long times for all these γ, with some deviations at shorter times in opposite directions for positive and negative γ. Concomitantly, as γ increases, the timescale over which $C(\tau)$ decays, scaled by the overall turnover time of $\overline{L}$ epochs, grows.

There is no sign of marginality in the exact solution for $\gamma = 0$, $\Sigma = 0$—therefore the Red Queen phase should be robust to small perturbations, as evinced by these simulations for $\gamma \neq 0$. The arguments of Sec. V D suggest that the robustness extends as well to large Σ, and numerical results suggest that it extends up to modest—although *not all*—γ: see Sec. VI for more discussion of this point. In contrast to the LV model, simulations of the linearized resource model suggest that the Red Queen phase obtains for all $\kappa < 1$.

APPENDIX M: FLUCTUATIONS IN STEADY STATE WITH $\Sigma = 0$

Figures 17(a) and 17(b) show that, in the Red Queen steady state for the linearized model, L fluctuates around a mean value $\overline{L}$ which is proportional to D, with fluctuations that scale with $\sqrt{D}$. This scaling is consistent with expectations from the finite fragility of the Red Queen steady state—with evolutionary perturbations to the community not causing large changes and fluctuations therefore down by a factor of $1/\sqrt{D}$. Figures 17(c) and 17(d) show the same statistics for Υ. We expect $\overline{\Upsilon}$ to be $O(1)$ but is anomalously small in simulations (since $\Sigma = 0$), and it is possible that the coefficient of its $O(1)$ part is zero, with a $1/\sqrt{D}$—or some other—correction showing up in numerics. For $\Sigma = 0$, we know that $\overline{\Upsilon} = 0$ in the limit $\kappa \to 0$, $D \to \infty$, $\kappa^2 D = Q^2$ (which reduces to the $\gamma = 0$ Lotka-Volterra model), as well as when $\kappa \to 1$, since in the latter case Υ becomes nonpositive and asymptotes to zero from below. It is therefore possible that $\overline{\Upsilon} = 0$ for intermediate κ as well, but our results are inconclusive. Note that fluctuations in Υ have size $O(1/\sqrt{D})$ as expected.

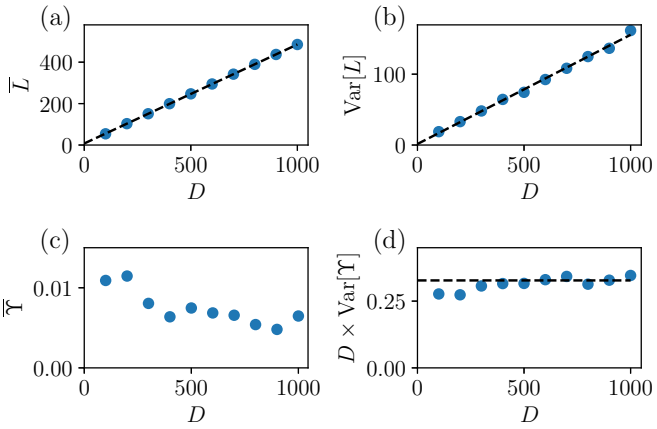


FIG. 17. Fluctuations of L and Υ in the Red Queen phase of the linearized model with $\kappa = 0.8$, $\rho = 0$, and $\Sigma = 0$. (a), (b) Both the mean and variance of L scale roughly linearly with D, so that fluctuations are smaller than the mean by a factor of $\sqrt{D}$. (c), (d) The mean Υ in steady state, expected to be $O(1)$, is rather small: unclear whether it goes to zero as $D \to \infty$. The variance of Υ scales, as expected, with $1/D$.

APPENDIX N: EVOLUTIONARY TRANSIENTS WITH $\Sigma > 0$

In Sec. V D we discussed the transient behavior of the evolution as the distribution of extant s_i pushes into the tail of $p(s)$. For a Gaussian distribution of the s_i, our arguments indicate that, during the transient,

$$\frac{\Upsilon}{\Sigma^2} \sim \left(\frac{T}{\Sigma^2}\right)^{1/3}. \tag{N1}$$

In Fig. 18 we find that this prediction agrees well with numerics for Υ as a function of evolutionary time, and the collapse of the curves scaled by Σ verifies the predictions. The curves collapse even for $\Sigma > \Sigma_\times$, beyond which the cavity analysis would naively predict that the mean strain lifetime diverges. The data indicate that, in agreement with observation in Fig. 11(c), for our choice of parameters, the simulations with $\Sigma \geqslant 4$ have not yet reached their steady state by the end of 5×10^6 epochs. The dynamics of L during the transient (not shown) are less well-fit by a simple power law since L varies over a smaller range and has a larger constant correction.

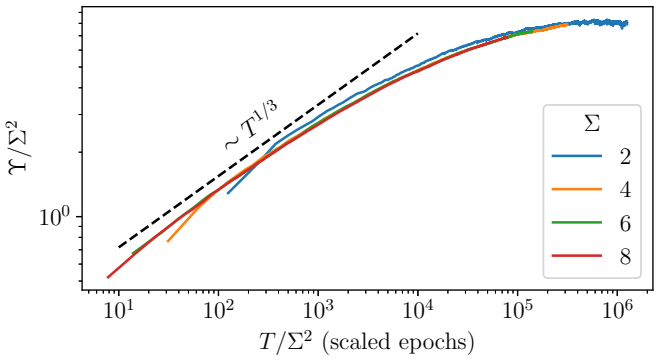


FIG. 18. Transient behavior of Υ during the evolutionary dynamics with large $\Sigma > \Sigma_\times \approx 1.9$ in the linearized model. Parameters are $D = 200$, $\kappa = 0.8$, and $\rho = 0$. Dashed black line has the predicted slope of $1/3$. Each simulation was run for 5×10^6 epochs.

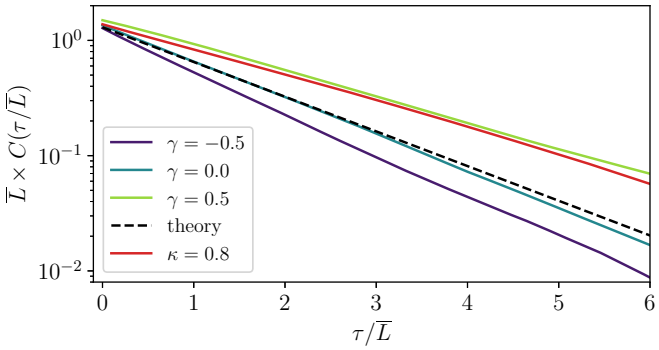


FIG. 19. In the LV model with $Q = 20$ and $\Sigma = 0$, $C(\tau/\overline{L})$ decays more slowly with increasing γ (normalization of τ by $\overline{L}$ eliminates the trivial scaling of the decay timescale with $\overline{L}$). Dashed black line is the simple-exponential prediction for $\gamma = 0$, with deviations from this result due to finite-size effects. Red line shows $C(\tau)$ for linearized model with $D = 400$, $\kappa = 0.8$.

APPENDIX O: CORRELATIONS AND LIFETIMES FOR NONZERO γ AND Σ

Here we examine the lifetime distribution, $p_{\mathcal{T}}(\mathcal{T})$ and the correlation function $C(\tau) = \sum_j \nu_j(T)\nu_j(T+\tau)$ in the LV model for various γ with $\Sigma = 0$, and for Gaussian $p(s)$ with various Σ, restricting to $\gamma = 0$. We find that in all cases—even when the cavity analysis suggests otherwise (Sec. V D)—a steady state appears to obtain in numerics, and the lifetime distribution decays sufficient rapidly that the relation $\langle\mathcal{T}\rangle = \overline{L}$ between the average strain lifetime and the average diversity holds.

1. Nonzero γ with $\Sigma = 0$

In Fig. 19 we show $C(\tau)$ for the $\Sigma = 0$ LV model, with time rescaled by the average strain lifetime, equal to $\overline{L}$. Defining rescaled time $\tilde{\tau} = \tau/\overline{L}$, we see that $C(\tilde{\tau})$ decays exponentially for large $\tilde{\tau}$, with the rate of this exponential depending on γ. Larger γ results in slower turnover and slower decay, presumably because with positive γ, invading strains are likely to increase the abundance of already-abundant strains, which therefore decay to extinction more slowly. Before the exponential regime it appears that for $\gamma < 0$, $C(\tilde{\tau})$ decays slightly faster for small $\tilde{\tau}$ than for large $\tilde{\tau}$, with the opposite trend for $\gamma > 0$. We also see exponential decay of the correlation function for the linearized model with $\kappa = 0.8$.

2. Nonzero Σ with $\gamma = 0$

In the LV model, due to the interactions being $O(1)$ [instead of $O(\sqrt{D})$ in the linearized model], the natural scale for the width of the extant s_i distribution at steady state is $O(1/Q) = O(1/\sqrt{L})$, rather than $O(1)$. We therefore consider $\Sigma = O(1/Q)$ in the following analysis.

For $\Sigma > 0$, subtleties arise for Gaussian $p(s)$ as discussed in Sec. V D. For Σ less than some crossover scale $\Sigma_\times$, $p_{\mathcal{T}}(\mathcal{T})$ is predicted to decay as a Σ-dependent *power law*, while

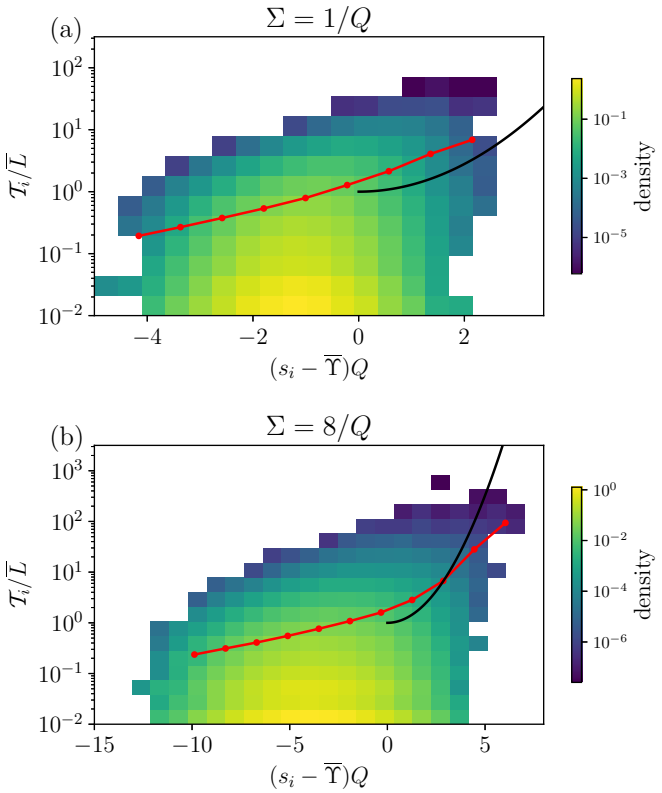


FIG. 20. Joint distribution of strain lifetimes $\mathcal{T}_i$ and general fitnesses s_i in the fully asymmetric LV model ($Q = 10$, yielding $\overline{L} \approx 70$). (a) Data for $\Sigma Q = 1$. Red line shows the mean lifetime $\langle \mathcal{T}|s\rangle$ conditioned on s within a given range. Black line shows naive theoretical expectation $\langle \mathcal{T}|s\rangle \sim \bar{L}\exp[(s-\overline{\Upsilon})^2/2\sigma_\zeta^2]$ (up to an unknown coefficient). Data show $\langle \mathcal{T}|s\rangle$ increasing more slowly than this expectation. (b) Same plot for $Q\Sigma = 8$, also showing empirical lifetime shorter than expectation from Gaussian steady state *Ansatz*.

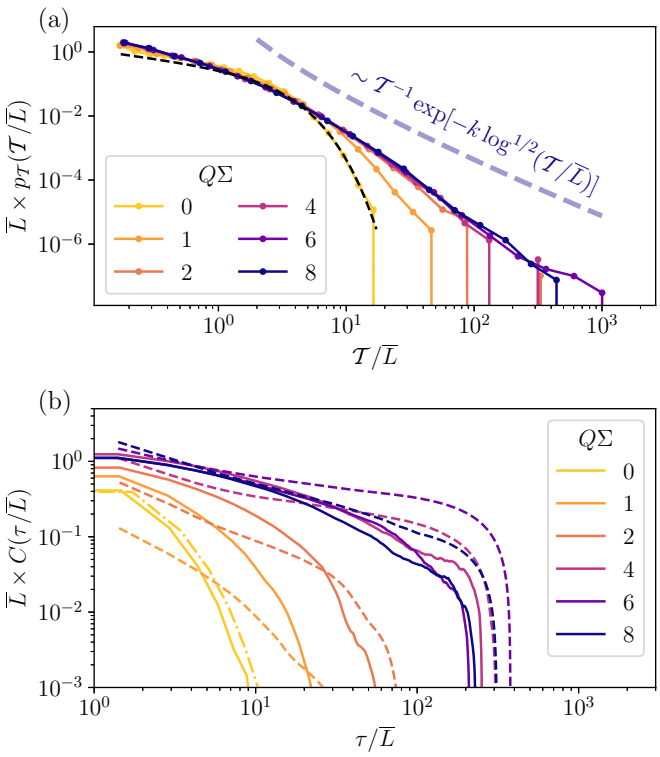


FIG. 21. (a) Distribution of lifetimes $\mathcal{T}_i$ for various values of Σ in the fully asymmetric LV model ($Q = 10$, yielding $\overline{L} \approx 70$). $p_{\mathcal{T}}(\mathcal{T})$ appears to decay sufficiently rapidly to yield finite $\langle \mathcal{T}\rangle$ even for large Σ. Thin dashed line shows theoretical prediction for $\Sigma = 0$ [Eq. (11)], which decays exponentially for large $\mathcal{T}$. Thick dashed line shows poor agreement with naive prediction (for $\Sigma = 8/Q$) for transient in $p_{\mathcal{T}}(\mathcal{T})$ of form shown with $k = \sqrt{2}\sigma_\zeta/\overline{\Delta} = \sqrt{2}\sigma_\zeta\overline{\Upsilon}/\Sigma^2 \cong 3.6$, and unknown multiplicative coefficient. (b) Correlation function $C(\tau)$ for various Σ. Yellow dash-dotted line shows theoretical result for $\Sigma = 0$ [Eq. (9)], with deviations from this prediction due to finite Q. Other dashed lines show predictions of $C(\tau)$ computed from $p_{\mathcal{T}}(\mathcal{T})$ using Eq. (13), with a similar number of strains averaged in $p_{\mathcal{T}}(\mathcal{T})$ as in $C(\tau)$. Note that time ranges in panels (a) and (b) are not aligned between panels.

for $\Sigma > \Sigma_\times$ [or for broader-than-Gaussian $p(s)$] one naively expects that large-s_i strains cause $\langle \mathcal{T}\rangle$ to diverge, disallowing a steady state. However, in contrast with this expectation, we find in numerics that a seemingly well-behaved steady state occurs even for large Σ.

To understand this behavior it is useful to examine deviations from the cavity analysis exhibited by our numerics. As discussed in Sec. V D, extinction for large $s - \overline{\Upsilon}$ requires a rare fluctuation of ζ_i away from zero (its mean) by an amount $s_i - \overline{\Upsilon}$. If the stochastic process $\zeta(T)$ has a *Gaussian* steady state even out into the rare tails of large negative ζ, the rate of extinction (and therefore inverse lifetime) will be roughly $\approx \exp[-(s_i - \overline{\Upsilon})^2/2C(0)]$, where $C(0) = \sigma_\zeta^2$ is the variance of the stationary state distribution of $\zeta(T)$. Figure 20 shows the joint distribution of $\mathcal{T}_i$ and $s_i - \overline{\Upsilon}$, and compares the mean lifetime for a given s, denoted $\langle \mathcal{T}|s\rangle$, to the predicted inverse-Gaussian-large dependence. The numerics show a slower increase of $\langle \mathcal{T}|s\rangle$ with s than suggested by a Gaussian steady state, indicating wider fluctuations of $\zeta(T)$ and shorter lifetimes.

These shorter-than-expected lifetimes can also be seen directly from $p_{\mathcal{T}}(\mathcal{T})$. Under the Gaussian cavity assumption, strains which last for time $\mathcal{T}$ will have typical general fitness $s \approx \overline{\Upsilon} + \sigma_\zeta\sqrt{2}\log^{1/2}(\mathcal{T}/\overline{L})$. If $p(s) \sim e^{-(s/\Sigma)^\psi/\psi}$, this yields an estimate for the large-$\mathcal{T}$ decay of $p_{\mathcal{T}}(\mathcal{T})$: for $\psi > 2$, we have $p_{\mathcal{T}}(\mathcal{T}) \sim \mathcal{T}^{-1}\exp[-k\log^{\psi/2}\mathcal{T}]$, while for Gaussian $p(s)$ we have, at asymptotically long times,

$$p_{\mathcal{T}}(\mathcal{T}) \sim \frac{1}{\mathcal{T}^{1+\beta}} \quad \text{with} \quad \beta = \sigma_\zeta^2/\Sigma^2, \tag{O1}$$

which yields a finite $\langle \mathcal{T}\rangle$ only for $\beta > 1$, i.e., $\Sigma < \Sigma_\times = \sigma_\zeta(\Sigma_\times) \cong 1.4/Q$. At intermediate times, especially for larger Σ, $p_{\mathcal{T}}(\mathcal{T})$ is expected to decay as $\approx \mathcal{T}^{-1}\exp[-k\log^{1/2}(\mathcal{T}/\overline{L})]$ with $k = \sqrt{2}\sigma_\zeta/\overline{\Delta}$, due to the roughly exponential decay of $p(s)$ (with slope $1/\overline{\Delta}$) near $\overline{\Upsilon}$.

In Fig. 21(a) $p_{\mathcal{T}}(\mathcal{T})$ is plotted on a log-log scale for Gaussian $p(s)$ with various Σ. As Σ goes from 0 to $8/Q$, σ_ζ increases modestly from $1.4/Q$ to $1.5/Q$, and $\overline{L}$ barely changes [Figs. 22(a) and 22(c)]. Even for $\Sigma = 1/Q$, which is less than $\Sigma_\times$, the decay of $p_{\mathcal{T}}(\mathcal{T})$ is considerably faster than the predicted $\beta + 1 \cong 3$. For $\Sigma \gg \sigma_\zeta$ the difference between numerics and cavity analysis is even larger: β appears to

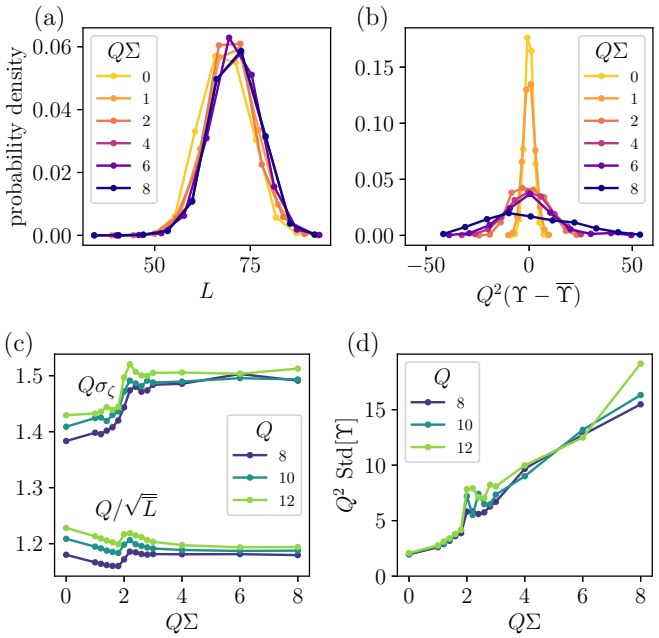


FIG. 22. Fluctuations in fully asymmetric LV model as a function of Σ. (a) Distributions of L showing that $\overline{L}$ and the fluctuations of L depend only weakly on Σ. (b) Distributions of fluctuations in Υ (whose mean scales as $Q\Sigma^2$) increase with Σ. The scale of these fluctuations is Q^2 at $\Sigma = 0$. $Q = 10$ in panels (a) and (b). (c) Scaled σ_ζ and scaled $1/\sqrt{L}$ as a function of scaled Σ (different sets of curves are labeled with colors indicating Q). The naive critical $\Sigma_\times$, defined by $\Sigma_\times = \sigma_\zeta(\Sigma_\times)$, is roughly $\Sigma_\times \approx 1.4/Q$ near which $1/\sqrt{L}$ and σ_ζ both exhibit a small increase. (d) Υ fluctuations increase dramatically as Σ passes through $\Sigma_\times$. For smaller Σ, Std[Υ], the standard deviation of Υ in steady state, scales as $1/Q^2$, while for larger Σ it appears to decrease more slowly with Q.

saturate at $\cong$1.7 in contrast with the predicted $\beta \cong 0.06$ for $\Sigma = 8/Q$, and $p_{\mathcal{T}}(\mathcal{T})$ is not well fit by the intermediate-$\mathcal{T}$ transient either [Fig. 21(a)]. Thus something is making the rare fluctuations to large negative ζ—assuming the basic cavity structure in terms of ζ is still valid—much larger than Gaussian.

Given the difference between naive expectations and numerics for $p_{\mathcal{T}}(\mathcal{T})$, we investigate whether these differences carry over to the correlation function, $C(\tau)$, which, as discussed in Sec. V D, is determined by strains that survive for at least time τ. $C(\tau)$ is plotted in Fig. 21(b) for Gaussian $p(s)$ with the same range of Σ between zero and $8/Q$. With the joint distribution of s and $\mathcal{T}$, and approximating the contribution to $C(\tau)$ of a strain with $\mathcal{T} > \tau$ by $(\mathcal{T} - \tau)\nu^2$, $C(\tau)$ can be estimated from the *observed* $p_{\mathcal{T}}(\mathcal{T})$ [Eq. (13)]. The result of this estimation is shown in Fig. 21(b) along with the empirically calculated $C(\tau)$, and we find rough agreement.

3. Reasons for breakdown of naive analysis

One possible source of the discrepancy between naive cavity analysis and numerics is the emergence of long-tailed memory kernels that increase the rate of extinction by allowing anomalously long-lasting negative ζ_i. Such dynamics could drive strains with large $s_i - \Upsilon$ extinct at a higher rate than the naive "barrier-crossing" estimate. But as long as the $\zeta(T)$ is a stationary Gaussian process, this cannot happen: the correlations between $\zeta(T)$ and the typically small ζ at the invasion time can only decrease the probability of a large deviation. A more viable scenario is that the large-$\overline{L}$ limit is subtle with large deviations from the Gaussian *Ansatz* of the $\zeta(T)$ statistics making extinction events of large $s - \overline{\Upsilon}$ strains much less unlikely, as seen in Fig. 20.

In general, one does not expect that the limits of large $\overline{L}$ and behavior in the tails of distributions will commute: the central limit theorem is only for the not-too-rare center of the distribution. Likewise, the long-time limit and large $\overline{L}$ limit need not commute. Concretely, if large fluctuations that drive extinctions are only exponentially rare in ζ, and hence lifetimes "only" exponentially long in $s - \overline{\Upsilon}$, $p_{\mathcal{T}}(\mathcal{T})$ would decay roughly as a power of $\mathcal{T}$ even for very large Σ where the distribution of extant s is close to exponential. This behavior is roughly consistent with the apparent power-law decay in this regime [Fig. 21(a)]. Exponential—rather than Gaussian—tails often occur, for example if the rare events that drive extinctions result from the sum of effects of a series of roughly independent events. Thus this putative resolution of the apparent paradox may well be viable.

Another effect that could cause extinctions is large fluctuations in Υ, particularly for exponential $p(s)$, which plays a key role in the behavior for large Σ. With exponential $p(s)$ (in contrast with $\psi > 1$) there is no special value of Υ in steady state. Instead there is a special value of Σ for which a steady state occurs, equal to $\overline{\Delta}$ in the steady state for $\psi > 1$. Since any Υ in an exponential tail is equivalent (conditioned on invasion of strains) there are likely to be at least diffusive—and hence unbounded—fluctuations in Υ. Upwards fluctuations in Υ would eventually drive extinctions of strains that had anomalously large $s - \Upsilon$ at their invasion time. Additionally, persistent large-s strains will themselves cause substantial increases in Υ, perhaps providing sufficient feedback for these to limit their own lifetimes.

The fluctuations in Υ about the steady state for Gaussian $p(s)$ can be estimated by noting that in the LV model, an invader changes Υ by an amount of order $L^{-3/2}$, the typical amount needed to keep $\sum_i \nu_i = 1$ for the random changes in ν_i caused by the invasion (Sec. V B). Combining these diffusive fluctuations with the $O(1/L^2\Sigma^2)$ linear deterministic push towards $\overline{\Upsilon}$ [which can be seen by expanding Eq. (16) for Υ near $\overline{\Upsilon}$] yields a prediction: $\mathrm{Var}(\Upsilon) \sim \Sigma^2/Q^2$. Simulations for modest $Q = 10$ in Fig. 22 show that while the distribution of L in steady state depends only weakly on Σ, the distribution of Υ in steady state appears consistent with this dependence on Σ. We expect that the diffusion coefficient of Υ would increase—and perhaps diverge—for large Σ, due to long-time correlations of the fluctuations in ζ and hence in Υ. To play a role in the extinctions of strains with large $s - \overline{\Upsilon}$, fluctuations of Υ would have to be larger than expected by a factor of $\sqrt{L} \sim Q$, and would likely need to have non-Gaussian tails as well. While we have not explored the tails, Figs. 22(b) and 22(d) indeed show increased fluctuations in Υ as a function of increasing Σ, with a sharp increase as Σ passes through $\Sigma_\times$ at which the naive analysis predicts a divergence of $\langle\mathcal{T}\rangle$. For larger Σ, the dependence of $Q^2\,\mathrm{Std}[\Upsilon]$ on Q becomes stronger, perhaps indicating a crossover to a regime with the fluctuations scaling as $1/Q$ and thus comparable to σ_ζ.

At this point it is unclear what kind of non-Gaussian fluctuations, or other subtleties of large but finite L, cause the persistence of a Red Queen steady state even when large s can occur. However, our numerics indicate that the eco-evolutionary steady state obtains even for large Σ with Gaussian $p(s)$. Further investigation of $\psi < 2$ is needed, but it is possible that a similar mechanism preserves the steady state there as well. It is worth noting that finite-$\overline{L} \approx 100$ corrections to the cavity analysis could be substantial for the range of parameters we investigated here: maybe such corrections enhance the robustness of the steady state for large Σ.

The above discussion has focused on the simplest case of the fully asymmetric ($\gamma = 0$) LV model. For $\gamma \neq 0$, or in the linearized model with κ independent of D, there are additional complications that could arise for $\Sigma > 0$ when the correlation function develops a long tail, which will induce long tails to both the memory kernels. It is unclear whether the persistent memory in $M(\tau)$ and $B(\tau)$ [Eqs. (L2) and (L3)], will change the behavior qualitatively and make the simple analysis break down even when it is still valid for $\gamma = 0$. However, the numerics show that the naive analyses are roughly correct for average properties of the evolved communities.

APPENDIX P: COMPARISON WITH RESULTS OF DE PIREY AND BUNIN

Recent work of de Pirey and Bunin [50] (dPB for short) has explored a Red Queen-like phase in a different context. dPB study *ecological dynamics* in a $\gamma = 0$ Lotka-Volterra model in the limit of small migration for each species, such that there is a clear separation between extant species which have substantial abundance, and species which would be extinct if not for the small migration. They find that if the diversity is high enough, the ecological dynamics are chaotic such that there are species which go nearly extinct before rising again to high abundance and rejoining the extant subset which is almost at a stable fixed point. Due to a separation of timescales arising from the very small migration rate (analogous to our separation of evolutionary and ecological timescales), the ecological dynamics consist of a series of almost-fixed points punctuated by changes that occur when a near-extinct species rises up to non-negligible abundance.

When a species transitions from near-extinct to extant, the process is similar to an independent invader joining the community, as the abundance of the invader can jump up to a nonzero value determined by its ecological interactions with the extant community. This entrance to the community causes a change in the biases of the other strains, which can drive low abundance strains out and produces a steady state in which low-abundance strains are depleted relative to the truncated Gaussian distribution of the assembled community. Indeed dPB find that their distribution at steady state is not well approximated by a truncated Gaussian, and the number of extant species at any given time is less than the total capacity of an assembled community.

However, in the model studied by dPB, the steady-state abundance distribution does not vanish at zero abundance, in contrast with our evolved communities. This is because when a species goes from extant to nearly extinct, its probability of rejoining the extant community due to an ecological fluctuation is a decreasing function of time as ecological fluctuations drive its abundance farther from the boundary between extant and near-extinct. Therefore, the boundary condition at zero abundance is not an absorbing one but rather a permeable boundary where species that hit zero abundance can be resurrected by a small change in the community and come back into the community soon after their extinction. dPB are able to write dynamical mean-field equations for the abundance dynamics and solve these equations numerically to get the steady-state abundance distribution. Our case turns out to be simpler in this respect since with evolutionary dynamics the extinction boundary is strictly absorbing which, fortuitously, makes the self-consistent correlations simple and enables an exact solution, with the equation satisfied by the abundance distribution reduced to a boundary-value problem [Eq. (K13)] that can be solved numerically (Fig. 5).

[1] S. Sunagawa, L. P. Coelho, S. Chaffron, J. R. Kultima, K. Labadie, G. Salazar, B. Djahanschiri, G. Zeller, D. R. Mende, A. Alberti *et al.*, Structure and function of the global ocean microbiome, Science **348**, 1261359 (2015).

[2] The Human Microbiome Project Consortium, A framework for human microbiome research, Nature (London) **486**, 215 (2012).

[3] N. Fierer, Embracing the unknown: Disentangling the complexities of the soil microbiome, Nat. Rev. Microbiol. **15**, 579 (2017).

[4] N. Kashtan, S. E. Roggensack, S. Rodrigue, J. W. Thompson, S. J. Biller, A. Coe, H. Ding, P. Marttinen, R. R. Malmstrom, R. Stocker *et al.*, Single-cell genomics reveals hundreds of coexisting subpopulations in wild prochlorococcus, Science **344**, 416 (2014).

[5] M. J. Rosen, M. Davison, D. Bhaya, and D. S. Fisher, Fine-scale diversity and extensive recombination in a quasisexual bacterial population occupying a broad niche, Science **348**, 1019 (2015).

[6] R. F. Rosenzweig, R. R. Sharp, D. S. Treves, and J. Adams, Microbial evolution in a simple unstructured environment: Genetic differentiation in *Escherichia Coli*, Genetics **137**, 903 (1994).

[7] M. L. Friesen, G. Saxer, M. Travisano, and M. Doebeli, Experimental evidence for sympatric ecological diversification due to frequency-dependent competition in *Escherichia Coli*, Evolution **58**, 245 (2004).

[8] R. Maharjan, S. Seeto, L. Notley-McRobb, and T. Ferenci, Clonal adaptive radiation in a constant environment, Science **313**, 514 (2006).

[9] B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, and M. M. Desai, The dynamics of molecular evolution over 60,000 generations, Nature (London) **551**, 45 (2017).

[10] R. E. Lenski, Experimental evolution and the dynamics of adaptation and genome evolution in microbial populations, ISME J. **11**, 2181 (2017).

[11] S. E. Finkel and R. Kolter, Evolution of microbial diversity during prolonged starvation, Proc. Natl. Acad. Sci. USA **96**, 4023 (1999).

[12] J. C. Ignacio-Espinoza, N. A. Ahlgren, and J. A. Fuhrman, Long-term stability and Red Queen-like strain dynamics in marine viruses, Nat. Microbiol. **5**, 265 (2020).
[13] M. F. Marston, F. J. Pierciey Jr, A. Shepard, G. Gearin, J. Qi, C. Yandava, S. C. Schuster, M. R. Henn, and J. B. H. Martiny, Rapid diversification of coevolving marine synechococcus and a virus, Proc. Natl. Acad. Sci. USA **109**, 4544 (2012).
[14] A. Buckling and P. B. Rainey, Antagonistic coevolution between a bacterium and a bacteriophage, Proc. R. Soc. London, Ser. B **269**, 931 (2002).
[15] L. Van Valen, A new evolutionary law, Evol. Theory **1**, 1 (1973).
[16] L. Carroll, *Through the Looking-Glass, and What Alice Found There* (Macmillan, London, 1871).
[17] J. A. G. M. de Visser and J. Krug, Empirical fitness landscapes and the predictability of evolution, Nat. Rev. Genet. **15**, 480 (2014).
[18] D. S. Fisher, Inevitability of Red Queen evolution driven by organismic complexity and simple feedback via environmental modification, bioRxiv 2021.09.26.461893; doi: https://doi.org/10.1101/2021.09.26.461893.
[19] A. Agarwala and D. S. Fisher, Adaptive walks on high-dimensional fitness landscapes and seascapes with distance-dependent statistics, Theor. Popul. Biol. **130**, 13 (2019).
[20] G. Bunin, Ecological communities with Lotka-Volterra dynamics, Phys. Rev. E **95**, 042414 (2017).
[21] M. Opper and S. Diederich, Phase transition and $1/f$ noise in a game dynamical model, Phys. Rev. Lett. **69**, 1616 (1992).
[22] E. Blumenthal, J. W. Rocks, and P. Mehta, Phase transition to chaos in complex ecosystems with nonreciprocal species-resource interactions, Phys. Rev. Lett. **132**, 127401 (2024).
[23] V. Ros, F. Roy, G. Biroli, G. Bunin, and A. M. Turner, Generalized Lotka-Volterra equations with random, nonreciprocal interactions: The typical number of equilibria, Phys. Rev. Lett. **130**, 257401 (2023).
[24] M. T. Pearce, A. Agarwala, and D. S. Fisher, Stabilization of extensive fine-scale diversity by ecologically driven spatiotemporal chaos, Proc. Natl. Acad. Sci. USA **117**, 14572 (2020).
[25] C. Xue and N. Goldenfeld, Coevolution maintains diversity in the stochastic "kill the winner" model, Phys. Rev. Lett. **119**, 268101 (2017).
[26] J. Huisman, A. M. Johansson, E. O. Folmer, and F. J. Weissing, Towards a solution of the plankton paradox: The importance of physiology and life history, Ecol. Lett. **4**, 408 (2001).
[27] L. Yan, R. A. Neher, and B. I. Shraiman, Phylodynamic theory of persistence, extinction and speciation of rapidly adapting pathogens, Elife **8**, e44205 (2019).
[28] R. A. Armstrong and R. McGehee, Competitive exclusion, Am. Nat. **115**, 151 (1980).
[29] N. Shoresh, M. Hegreness, and R. Kishony, Evolution exacerbates the paradox of the plankton, Proc. Natl. Acad. Sci. USA **105**, 12365 (2008).
[30] P. Schippers, A. M. Verschoor, M. Vos, and W. M. Mooij, Does "supersaturated coexistence" resolve the "paradox of the plankton"? Ecol. Lett. **4**, 404 (2001).
[31] M. Doebeli and I. Ispolatov, Complexity and diversity, Science **328**, 494 (2010).
[32] M. Tikhonov and R. Monasson, Innovation rather than improvement: A solvable high-dimensional model highlights the limitations of scalar fitness, J. Stat. Phys. **172**, 74 (2018).
[33] C. A. Serván and S. Allesina, Tractable models of ecological assembly, Ecol. Lett. **24**, 1029 (2021).
[34] J. A. Capitán, J. A. Cuesta, and J. Bascompte, Statistical mechanics of ecosystem assembly, Phys. Rev. Lett. **103**, 168101 (2009).
[35] J. McEnany and B. H. Good, Predicting the first steps of evolution in randomly assembled communities, Nat. Commun. **15**, 8495 (2024).
[36] A. Altieri, F. Roy, C. Cammarota, and G. Biroli, Properties of equilibria and glassy phases of the random Lotka-Volterra model with demographic noise, Phys. Rev. Lett. **126**, 258301 (2021).
[37] M. Tikhonov and R. Monasson, Collective phase in resource competition in a highly diverse ecosystem, Phys. Rev. Lett. **118**, 048103 (2017).
[38] W. Cui, R. Marsland III, and P. Mehta, Effect of resource dynamics on species packing in diverse ecosystems, Phys. Rev. Lett. **125**, 048101 (2020).
[39] G. Biroli, G. Bunin, and C. Cammarota, Marginally stable equilibria in critical ecosystems, New J. Phys. **20**, 083051 (2018).
[40] R. Marsland III, W. Cui, and P. Mehta, The minimum environmental perturbation principle: A new perspective on niche theory, Am. Nat. **196**, 291 (2020).
[41] J. A. Bonachela, M. T. Wortel, and N. C. Stenseth, Eco-evolutionary Red Queen dynamics regulate biodiversity in a metabolite-driven microbial system, Sci. Rep. **7**, 17655 (2017).
[42] R. Law, P. Marrow, and U. Dieckmann, On evolution under asymmetric competition, Evol. Ecol. **11**, 485 (1997).
[43] J. M. Nordbotten and N. C. Stenseth, Asymmetric ecological conditions favor red-queen type of continued evolution over stasis, Proc. Natl. Acad. Sci. USA **113**, 1847 (2016).
[44] R. M. May, Will a large complex system be stable? Nature (London) **238**, 413 (1972).
[45] A. Mahadevan, M. T. Pearce, and D. S. Fisher, Spatiotemporal ecological chaos enables gradual evolutionary diversification without niches or tradeoffs, eLife **12**, e82734 (2023).
[46] A. Posfai, T. Taillefumier, and N. S. Wingreen, Metabolic trade-offs promote diversity in a model ecosystem, Phys. Rev. Lett. **118**, 028103 (2017).
[47] R. Caetano, Y. Ispolatov, and M. Doebeli, Evolution of diversity in metabolic strategies, eLife **10**, e67764 (2021).
[48] B. H. Good, S. Martis, and O. Hallatschek, Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources, Proc. Natl. Acad. Sci. USA **115**, E10407 (2018).
[49] C. Cockrell, J. D. O'Sullivan, J. C. D. Terry, E. C. Nwankwo, K. Trachenko, and A. G. Rossberg, Self-organization of ecosystems to exclude half of all potential invaders, Phys. Rev. Res. **6**, 013093 (2024).
[50] T. Arnoulx de Pirey and G. Bunin, Many-species ecological fluctuations as a jump process from the brink of extinction, Phys. Rev. X **14**, 011037 (2024).
[51] A. J. Daly, J. M. Baetens, and B. De Baets, Ecological diversity: Measuring the unmeasurable, Mathematics **6**, 119 (2018).
[52] F. Vrins, Sampling the multivariate standard normal distribution under a weighted sum constraint, Risks **6**, 64 (2018).
[53] F. Roy, G. Biroli, G. Bunin, and C. Cammarota, Numerical implementation of dynamical mean field theory for disordered systems: Application to the lotka–volterra model of ecosystems, J. Phys. A: Math. Theor. **52**, 484001 (2019).

[54] L. F. Cugliandolo, J. Kurchan, P. Le Doussal, and L. Peliti, Glassy behaviour in disordered systems with nonrelaxational dynamics, Phys. Rev. Lett. **78**, 350 (1997).

[55] M. Bonsma-Fisher and S. Goyal, Dynamics of immune memory and learning in bacterial communities, eLife **12**, e81692 (2023).

[56] J. Marchi, M. Lässig, A. M. Walczak, and T. Mora, Antigenic waves of virus-immune coevolution, Proc. Natl. Acad. Sci. USA **118**, e2103398118 (2021).

[57] A. Goyal, L. S. Bittleston, G. E. Leventhal, L. Lu, and O. X. Cordero, Interactions between strains govern the eco-evolutionary dynamics of microbial communities, eLife **11**, e74987 (2022).

[58] T. Yoshida, S. P. Ellner, L. E. Jones, B. J. M. Bohannan, R. E. Lenski, and N. G. Hairston, Jr., Cryptic population dynamics: rapid evolution masks trophic interactions, PLoS Biol. **5**, e235 (2007).

[59] M. Valles-Colomer, A. Blanco-Míguez, P. Manghi, F. Asnicar, L. Dubois, D. Golzato, F. Armanini, F. Cumbo, K. D. Huang, S. Manara *et al.*, The person-to-person transmission landscape of the gut and oral microbiomes, Nature (London) **614**, 125 (2023).

[60] W. Zhou, M. Spoto, R. Hardy, C. Guan, E. Fleming, P. J. Larson, J. S. Brown, and J. Oh, Host-specific evolutionary and transmission dynamics shape the functional diversification of staphylococcus epidermidis in human skin, Cell **180**, 454 (2020).

[61] A. M. Martin-Platero, B. Cleary, K. Kauffman, S. P. Preheim, D. J. McGillicuddy, E. J. Alm, and M. F. Polz, High resolution time series reveals cohesive but short-lived communities in coastal plankton, Nat. Commun. **9**, 266 (2018).

[62] S. Diederich and M. Opper, Replicators with random interactions: A solvable model, Phys. Rev. A **39**, 4333 (1989).

[63] L. M. Ricciardi and S. Sato, First-passage-time density and moments of the Ornstein-Uhlenbeck process, J. Appl. Probab. **25**, 43 (1988).